

REPLY

Collective Obligations and the Institutional Critique of Effective Altruism: A Reply to Alexander Dietz

Brian Berkey

University of Pennsylvania

Corresponding author. Email: bberkey@wharton.upenn.edu

(Received 30 September 2018; revised 7 March 2019; accepted 7 March 2019)

In a recent article in this journal, Alexander Dietz argues that what I have called the ‘institutional critique of effective altruism’ is best understood as grounded in the claim that ‘EA relies on an overly individualistic approach to ethics, neglecting the importance of our collective obligations’. In this reply, I argue that Dietz’s view does not represent a plausible interpretation of the institutional critiques offered by others, primarily because, unlike Dietz, they appear to believe that their critiques provide reasons to reject the EA view about the content of our individual obligations. I also argue that EA’s identity as a social movement provides grounds for denying Dietz’s claim that it is objectionably incomplete.

I. Introduction

In a recent article in this journal, Alexander Dietz argues that what I have called the ‘institutional critique of effective altruism’¹ is best understood as grounded in the claim that ‘EA relies on an overly individualistic approach to ethics, neglecting the importance of our collective obligations’.² This neglect does not, on Dietz’s view, give us reason to think that EA’s central commitments, including the commitment to thinking that individuals should aim to do the most good they can, are mistaken. Instead, he claims that we should think that EA is ‘incomplete’, since there are reasons, grounded in EA’s commitment to doing the most good possible, to think that there are obligations to promote the good that are possessed by collectives as such.³ On this view, the charge that EA has wrongly neglected the importance of efforts to bring about large-scale, global institutional change is, in at least one sense, correct, since EAs should acknowledge that collectives that are capable of advancing efforts to bring about such change can be obligated to do so.

I have three central aims in this reply. First, in section II, I will attempt to clarify the relationship between the view defended by Dietz and the critiques of EA that I have

¹Brian Berkey, ‘The Institutional Critique of Effective Altruism’, *Utilitas* 30 (2018), pp. 143–71.

²Alexander Dietz, ‘Effective Altruism and Collective Obligations’, *Utilitas* 31 (2019), pp. 106–15, at 107.

³Dietz, ‘Effective Altruism and Collective Obligations’, pp. 107, 112–13. See also Dietz, ‘What We Together Ought to Do’, *Ethics* 126 (2016), pp. 955–82.

previously discussed under the heading of the ‘institutional critique’. I will argue that Dietz’s view does not represent a plausible interpretation of the critiques offered by others, primarily because, unlike Dietz, they appear to believe that their critiques provide reasons to reject the EA view about the content of our individual obligations. Next, in section III, I will consider whether Dietz’s assessment of EA as incomplete should be accepted. I will argue that EA’s identity as a social movement grounded in a limited set of core philosophical commitments gives members reason to deny that acceptance of the controversial view that collectives as such can be the bearers of obligations is among the core commitments. If this is correct, then Dietz’s claim that EA is incomplete should be rejected. I will conclude, in section IV, by considering how EAs should think about collective action and efforts to bring about global institutional change. I will claim that although Dietz is correct that EAs have reasons, grounded in their own central commitments, to be concerned that collectives, in addition to individuals, do more good rather than less, this does not require that they adopt any particular view about whether collectives as such can be the bearers of obligations.

II. Dietz’s view and the institutional critique of EA

Dietz’s argument for the view that EAs should accept that collectives, in addition to individuals, have obligations to promote the good relies on a discussion of cases with a particular structure. The central feature of the cases is that it is possible for each member of a collective to do the most good that she can, individually, despite the fact that the collective could have done more good had the individual members acted differently.⁴ To illustrate, consider a variant of a case that I discussed in my earlier article:⁵

*Drowning Children**: Five children are in danger of drowning. All five can be saved if, and only if, A, B and C work together to paddle a nearby canoe to the children. If fewer than three people attempt to use the canoe to rescue the children, they will arrive too late to save any of them. Each of A, B and C is also near enough to one life preserver to run to it in time to save one of the children. All three run to the life preservers, with the result that three of the five children are saved.

In this case, each individual has done as much good as possible, given what the others in fact do. Had any one of them run to the canoe rather than to the life preserver, one fewer child would have been saved. Nevertheless, it seems clear that something has gone wrong, since there was an alternative course of action available to the group that would have resulted in all five children being saved. Intuitively, we want to say that there is some sense in which all of the children ought to have been saved. A view about the obligations that exist in cases of this kind that says only that each individual should do as much good as she can, however, appears unable to capture this intuition, since it is possible for each individual to satisfy this obligation even if the group acts suboptimally.

Dietz claims that we can accommodate the intuition that all of the children ought to have been saved by holding that groups as such can be the bearers of obligations.⁶ On this view, the collective consisting of A, B and C was obligated to use the canoe to rescue

⁴Dietz, ‘Effective Altruism and Collective Obligations’, p. 109; Dietz, ‘What We Together Ought to Do’, pp. 960–3.

⁵Berkey, ‘The Institutional Critique of Effective Altruism’, pp. 155–6.

⁶Dietz, ‘Effective Altruism and Collective Obligations’, p. 111.

the children, and failed to satisfy this obligation. Crucially, for Dietz, this can be true even if it is also true that none of the individual members of the collective failed to do what she ought to have done.⁷ Given that A and B in fact ran to the life preservers, it seems plausible that C acted as she ought to have by doing so as well (and likewise for A and B). Endorsing collective obligations to do the most good in cases of this kind does not, according to Dietz, require that we reject the view that individuals also ought to do the most good that they can. This is why he claims that his version of the institutional critique implies only that EA is incomplete, and not that it is mistaken about what individuals ought to do.

Dietz's claim that the best way to capture the intuition that all of the children ought to have been saved is to endorse the view that collectives as such can be the bearers of obligations is quite plausible. It does not, however, seem to me to represent a plausible interpretation of what others who have offered versions of the institutional critique seem to find most objectionable about EA. While it seems clear that at least some of them would endorse Dietz's view that collectives as such can be the bearers of obligations, there are strong reasons to think that their primary concern is with EA's account of our individual obligations with respect to, for example, global poverty alleviation, and that insofar as they take it to be important that we think about our collective obligations, this is at least in part because doing so should, they think, lead us to reject the view that, as individuals, we are obligated to do the most good we can. Instead, they seem to think that individuals are at least permitted, and perhaps required, to devote at least some effort to promoting large-scale institutional change, regardless of whether there are strong reasons to think that this is what would do the most good.

Judith Lichtenberg, for example, responds to Will MacAskill's claim that, as individuals, we ought to donate to charitable organizations that will do the most good, rather than to organizations that happen to benefit people we know or causes to which we are especially attached,⁸ by claiming that 'to most of us it seems neither reasonable nor desirable to expect people to remain untouched by the particular individuals they know and the particular causes that affect them'.⁹ She also claims that we should conceive of duties to alleviate global poverty as 'belonging to collectives *rather than individuals*', and one of the central reasons that she offers for accepting this view is that 'such a shift would reduce the demands on individuals'.¹⁰ One of her main concerns about EA, then, is that its account of our individual obligations is objectionably demanding. Her claim that we ought to conceive of duties to alleviate poverty as fundamentally collective, then, is motivated at least in part by the thought that a view of this kind can accommodate the intuition that an individual obligation to do the most good is objectionably demanding. Rather than endorsing the kinds of individual obligations that EAs do, she suggests that individuals are permitted to give to a wide range of causes, from those that are clearly quite morally important, even if not where funds will do the most good, such as 'domestic poverty relief [and] racial justice', to those whose moral significance in comparison with issues like global poverty is at best questionable, such as

⁷Dietz, 'Effective Altruism and Collective Obligations', pp. 112–13.

⁸William MacAskill, *Doing Good Better: How Effective Altruism Can Help You Make a Difference* (New York, 2015), ch. 2.

⁹Judith Lichtenberg, 'Peter Singer's Extremely Altruistic Heirs', *New Republic* (November 2015), <<https://newrepublic.com/article/124690/peter-singers-extremely-altruistic-heirs>> [accessed 25 September 2018].

¹⁰Judith Lichtenberg, *Distant Strangers: Ethics, Psychology, and Global Poverty* (New York: 2014), p. 68, italics added.

‘religious organizations ... the opera, [or] your alma mater’.¹¹ Clearly, then, Lichtenberg, unlike Dietz, thinks that an important reason for endorsing collective obligations to address issues such as global poverty is that this commitment will affect what we should think about our individual obligations, and in particular will allow us to endorse a less demanding account of those obligations.

Like Lichtenberg, Amia Srinivasan suggests that it is at least permissible for individuals to devote time and resources to charitable efforts that they happen to care particularly about, even if these efforts do much less good than others.¹² She also thinks, then, that the EA view that, as individuals, we are obligated to do the most good is objectionable because it implies that we are obligated to forego supporting less effective efforts. She also suggests that the EA view that individuals should decide where to direct their efforts to improve the world on the basis of evidence about the relative effectiveness of different alternatives, and estimates of the value and probability of success in one’s efforts, is problematic because of the difficulty of assessing the value of radical political transformations and individual efforts to contribute to them:

What’s the expected marginal value of becoming an anti-capitalist revolutionary? To answer that you’d need to put a value and probability measure on achieving an unrecognisably different world – even, perhaps, on our becoming unrecognisably different sorts of people.¹³

It seems clear that Srinivasan thinks that becoming an anti-capitalist revolutionary is at least a permissible way that an individual might attempt to improve the world, even if there is no reason to think that it is among the options with the highest expected value. This, along with her criticism that MacAskill, and EA more generally, ‘does not address the deep sources of global misery – international trade and finance, debt, nationalism, imperialism, racial and gender-based subordination, war, environmental degradation, corruption, exploitation of labour – or the forces that ensure its reproduction’¹⁴ together suggest that she thinks that individuals have strong reasons to direct their efforts at addressing the large-scale institutional issues to which she refers rather than doing the other kinds of things that EAs tend to advocate, at least when they conflict.

Finally, Lisa Herzog, in describing how we should think about morality in light of her version of the institutional critique, says that:

morality is not about picking and choosing charities from an armchair; it’s about trying to become a force for change in daily life, and supporting whatever cause we can contribute to actively, passionately, and in ways that can create institutions and practices in line with our moral values and ideals.¹⁵

There is a lot that could be unpacked in this claim, but the important point for my purposes is that it seems clear that, like Lichtenberg and Srinivasan, Herzog takes her

¹¹Lichtenberg, ‘Peter Singer’s Extremely Altruistic Heirs’.

¹²Amia Srinivasan, ‘Stop the Robot Apocalypse’, *London Review of Books* 37 (2015), pp. 3–6, <<http://www.lrb.co.uk/v37/n18/amia-srinivasan/stop-the-robot-apocalypse>> [accessed 25 September 2018].

¹³Srinivasan, ‘Stop the Robot Apocalypse’.

¹⁴Srinivasan, ‘Stop the Robot Apocalypse’.

¹⁵Lisa Herzog, ‘Can “Effective Altruism” Really Change the World?’, *OpenDemocracy.net* (February 2016), <<https://www.opendemocracy.net/transformation/lisa-herzog/can-effective-altruism-really-change-world#>> [accessed 25 September 2018].

critique to have implications for how individuals should be directing their efforts in the actual world – specifically, they should be aiming at creating better institutions and practices, rather than (or in addition to, where there is no conflict) donating to effective charitable efforts.

Lichtenberg's, Srinivasan's and Herzog's discussions are, it seems to me, representative of the criticisms of EA that can plausibly be included under the 'institutional critique' label. And one thing that seems common to them is the view that the critique bears not just on abstract philosophical questions about whether collectives as such can be the bearers of obligations, and whether, if they can, this can help us capture intuitions about cases like *Drowning Children*^{*}, but also on questions about what individuals (in addition to collectives) ought to be doing in response to issues such as global poverty. Because of this, it seems to me that, whatever its philosophical merits, Dietz's view, according to which EA's account of our individual obligations is (or at least might be) correct, cannot be thought of as offering an interpretation of the institutional critique of EA, as it has been developed by others, that is faithful to their central concerns.

It is important to note that Dietz does not claim that his view is consistent with those of critics of EA, but only that it constitutes a version of the institutional critique that can avoid the objections to it that I offer in my earlier article.¹⁶ He does not, then, aim to defend the institutional critique as it has been developed by others, but instead to offer a view that can capture what he takes to be at least one of the central intuitions motivating proponents of the critique, namely that one thing that ought, in some relevant sense, to be done in response to global poverty and injustice is reforming unjust global political and economic institutions. My argument that his view does not represent a central concern of other critics, namely that EA's account of our individual obligations is mistaken, then, is not intended to suggest that he misunderstands their concerns or mischaracterizes their views, but instead to highlight the central differences between his criticism and those of others, which is necessary in order to assess his claim that EA, as it has been developed thus far, requires supplementation by the claim that collectives as such can have obligations to do the most good possible.

III. Is EA incomplete?

Dietz suggests that his argument shows that EA is incomplete, even if its claims about what individuals ought to do are correct. It seems to me, however, that my discussion of the relationship between his argument and those of others who have offered institutional critiques, along with a plausible view about the nature of EA as a 'philosophy and social movement',¹⁷ casts some doubt on this claim.

Critics like Lichtenberg, Srinivasan and Herzog take as their primary target the EA view about what individuals should be doing to try to improve the world. They think that EA is objectionable because, in some sense or other, it does not encourage individuals to focus as much of their effort as they should on attempting to bring about large-scale institutional change. This line of criticism treats EA as offering a view about how individuals ought to decide what to do given the actual state of the world, including facts about what others are likely to do. In other words, it treats EA as offering normative prescriptions for individuals at least primarily at the level of non-ideal

¹⁶Dietz, 'Effective Altruism and Collective Obligations', p. 107.

¹⁷Peter Singer, *The Most Good You Can Do: How Effective Altruism is Changing Ideas About Living Ethically* (New Haven, 2015), pp. 4–5.

theory. Dietz's argument, on the other hand, assumes that the EA view about what individuals ought to do in non-ideal conditions like those in the actual world is (or at least might be) correct, and goes on to suggest that this view requires supplementation by a view about what collectives as such are obligated to do. And while his view about what collectives as such are obligated to do can be thought to have implications for a wide range of cases at both the ideal and non-ideal theory levels, since it can in principle be applied to any collective that can act in any circumstances in which it might find itself, in order to capture what he takes to be the central valuable insight of the institutional critique, namely that there is a requirement, which EA fails to capture, to bring about the best possible large-scale institutional changes, the view must include a particular commitment at the level of ideal theory. Specifically, it must include a commitment to the view that the collective consisting of all of humanity is obligated to implement a fully just set of global institutions. This is because any failure of compliance by an agent, individual or collective, that is involved in a case of the relevant kind (in which we have an intuition that the optimal outcome, such as a fully just set of global institutions, ought to be brought about) will make it the case that a (larger) collective has failed to produce the required outcome.¹⁸ Dietz's argument, then, suggests that the incompleteness of EA consists, at least in part, in a failure to accept a claim at the level of ideal theory that, on his view, has no bearing on what individuals ought to do in non-ideal conditions.

Even if Dietz is correct that collectives as such can be the bearers of obligations to do the most good possible, however, whether this gives us reason to conclude that EA, as it has been developed thus far, is incomplete depends on what EA is and should aim to be. If EA is intended to be, or should be regarded as, a comprehensive moral theory, or a complete normative outlook, then Dietz's argument would, at least if it succeeds, show that EA is incomplete. It seems to me, however, that EA should not be thought of as, or as intended to be, a comprehensive moral theory or a complete normative outlook. Instead, it is, as Peter Singer suggests,¹⁹ a social movement structured around a set of core philosophical commitments. These core commitments, though they certainly include contested normative claims, do not amount to a complete normative outlook. And the fact that it is not a complete normative outlook allows EA, as a movement, to avoid taking positions on a range of controversial questions. This is important, since it allows EA to appeal to people with a range of views who might nonetheless be persuaded by arguments for the core commitments.

While there are challenging questions about exactly how the core commitments should be understood, it seems clear that EA is fundamentally a view about how individuals ought to decide what to do to try to improve the world in conditions like those in the actual world. This is, at least for the most part, how critics such as Lichtenberg, Srinivasan, and Herzog have treated it – they have offered arguments that aim to challenge the EA view about how individuals should direct their efforts to improve the world. But if EA's core commitments are, and ought to be, limited to issues related to individual obligations at the level of non-ideal theory, then Dietz's claim that it should be viewed as incomplete is in the relevant sense mistaken, even if his argument for the view that collectives as such can be the bearers of obligations otherwise succeeds.

To illustrate this point further, we can begin by imagining a social movement, much like EA, that is structured around the view represented in Singer's 'Famine, Affluence,

¹⁸Dietz, 'Effective Altruism and Collective Obligations', p. 113.

¹⁹Singer, *The Most Good You Can Do*, ch. 1.

and Morality'.²⁰ This movement, and its core commitments, might be challenged in ways that are familiar from the challenges to Singer's argument. It might be argued, for example, that the movement's view about individuals' obligations is objectionably demanding, or that it is mistaken because there are reasons for individuals to focus more of their efforts on promoting large-scale institutional change that the view cannot properly take into account. These kinds of concerns are the focus of most of the institutional critiques of EA – they treat EA as, fundamentally, a view about our individual obligations in conditions like ours, and challenge it on those terms.

On the other hand, a challenge to this movement on the grounds that it lacks, for example, a comprehensive theory of global justice would seem to miss its mark, for the same reason that a challenge to Singer's philosophical argument according to which it is objectionably incomplete in virtue of lacking such a theory would miss its mark. Because the argument aims to offer grounds for a normative conclusion that do not involve or require any very specific commitments regarding the correct comprehensive theory of global justice, objecting that the argument lacks such a theory fails to engage with it on its own terms. Of course there is a sense in which the view that the argument defends *is* incomplete – it does not constitute a comprehensive normative outlook – but this is not a compelling objection if the argument does not, and need not, aim to offer such a comprehensive outlook. Similarly, it is not a compelling objection to a social movement that it lacks a commitment to a comprehensive theory of global justice, if there are good reasons for it to avoid such a commitment.

Since there are good reasons for EA, as a social movement, to avoid adopting controversial philosophical positions (beyond those that constitute its core commitments) where it can, it is, it seems to me, a mistake to think that it is objectionably incomplete in virtue of failing to accept, among its core commitments, the view that collectives as such can be bearers of obligations.

IV. EA and collective action

Even if I am correct that EAs need not accept among their core commitments that collectives as such can be the bearers of obligations, Dietz has identified an important concern that EAs have reason to take seriously. When, as in *Drowning Children**, a group of people can do more good by cooperating than could be done if each individual acts alone to try to promote the good, EAs have reason to endorse and promote cooperation.

This concern is relevant not only in thinking about questions at the level of ideal theory, but also in thinking about what a range of groups can do in non-ideal conditions in which other agents, both individual and collective, are failing to comply with their obligations. As Dietz points out, we can, for example, think about what the EA movement as a collective might do to try to improve the world, or what a particular EA-aligned group might do.²¹ It could be the case, for example, that EAs pooling all of their charitable resources and directing them towards an effort to bring about certain institutional reforms would have greater expected value than would directing the same resources to EA-endorsed charities. If this is the case, then it seems clear that EAs should endorse pooling the resources.

Does this give us reason to think that Dietz is, in fact, correct that EAs have reason to accept the view that collectives as such can be bearers of obligations to do the most good

²⁰Peter Singer, 'Famine, Affluence, and Morality', *Philosophy & Public Affairs* 1 (1972), pp. 229–43.

²¹Dietz, 'Effective Altruism and Collective Obligations', p. 111.

possible? Must members of the EA community take the community itself to be the bearer of an obligation to pool resources and promote institutional reform in cases in which that would do more good than relatively uncoordinated charitable donations?

Since EA has reasons to avoid adopting philosophically controversial positions where it can do so without conflict with its core commitments, whether it must endorse Dietz's view depends on whether there are alternative, plausible ways of making sense of the need to endorse and promote cooperation in cases in which it is necessary for an EA group, or the movement as a whole, to do the most good. And it seems to me that there is at least one alternative type of view that shares with Dietz's compatibility with the commitment to the EA view that individuals should aim to do the most good that they can.

Views of this type are defended by Donald Regan and Robert Goodin. On Regan's version, individuals are obligated to cooperate, and to stand willing to cooperate, with anyone else who is willing to cooperate in order to produce the best possible outcome.²² On Goodin's version, each individual is obligated to contribute to a cooperative effort that would bring about a good outcome on the condition that enough others are willing to contribute if one does.²³ These views both require EAs to be willing, for example, to pool resources with other EAs in order to support institutional change efforts when this will do more expected good than alternative uses of the resources. More generally, they require individuals both to stand ready to contribute to cooperative efforts with others who are also willing, and to contribute when doing so will in fact bring about a better outcome, but not when it will not. In addition, accepting either of them does not commit one to accepting that collectives as such can be the bearers of obligations.²⁴

As a philosophical matter, these alternative views may be less plausible than Dietz's view. But given EA's identity as a social movement structured around core philosophical commitments regarding our individual obligations in non-ideal conditions, there is no reason that it must take a position one way or the other, and good reasons for it not to do so. Because of this, Dietz's claim that EA is incomplete should be rejected.²⁵

²²Donald Regan, *Utilitarianism and Cooperation* (Oxford, 1980), ch. 8.

²³Robert Goodin, 'Excused by the Unwillingness of Others', *Analysis* 72 (2012), pp. 18–24.

²⁴Dietz offers an objection to views of this kind that is, in my view, at least fairly compelling ('What We Together Ought to Do', p. 962). But because his aim in the article to which I am responding is to argue that EAs, qua EAs, ought to accept the view that collectives as such can be the bearers of obligations, what he would need to show in order for the argument to succeed is that alternative views such as Regan's and Goodin's cannot account for the requirement that EAs stand willing to cooperate, and in particular to pool resources, with others who are committed to contributing to bringing about the most good possible. Because their views *can* account for that requirement, there is no reason that EA as a movement needs to take a position on the success or failure of Dietz's objection to them.

²⁵I am grateful to the Edmond J. Safra Center for Ethics at Harvard University and the Berggruen Institute for their support of this work.