

Effective Animal Advocacy

Jeff Sebo
New York University

“[A] morality which never shocks anybody dwindles into etiquette. ... All the same, ideals which nobody can translate into action are wasted.”

- Mary Midgley, *Animals and why they matter*, p. 9

1. Introduction

Imagine that you are a doctor, volunteering your time to save people in the aftermath of a natural disaster. There are many more people who need help than you have ability to help. How should you approach your work? Many people find it natural to say that you should triage. That is, you should spend your resources as effectively and efficiently as possible, so that you can do the most good possible. If you have to choose between treating two people, one of whom has a major injury and the other of whom has a minor injury, then you should prioritize the person with the major injury, all else being equal. Similarly, if you have to choose between treating two people, one of whom requires relatively few scarce resources and the other of whom requires relatively many scarce resources, then you should prioritize the person who requires relatively few scarce resources, all else being equal.

Granted, it might seem callous to prioritize lives in this kind of way. But in a state of emergency, we naturally understand that triage is not an expression of callousness, but is rather an expression of informed, rational, impartial benevolence. If we want to save the most lives possible, or relieve the most suffering possible, then we need to think strategically about how best to use our limited resources so that we can achieve this aim.

However, as Will MacAskill (2015), Peter Singer (2015), and many others have noted, many people seem to forget this point when it comes to addressing other, more chronic problems.¹ For example, when we make choices about advocacy and philanthropy, many of us do what feels personally meaningful rather than think strategically

¹ MacAskill (2015) motivates the idea of effective altruism with a similar example (pp. 29-32). This introduction, as well as much of the rest of this piece, is indebted to his work.

about how to save the most lives possible or relieve the most suffering possible. Indeed, many of us spend more time researching which appliances to buy than which organizations to support with our time, energy, and money.

This is where effective altruism (EA) comes in. EA is a social movement that aims to use evidence and reason to do the most good possible. For an effective altruist, the world is in a chronic state of emergency, with many more individuals who need help than we have the ability to help. So, we should triage: We should spend our resources as effectively and efficiently as possible, so that we can do the most good possible. If we have to choose between supporting two organizations, one of which would save or spare more lives than the other as a result of our support, then we should prioritize the organization that would save or spare more lives, all else being equal.

Effective animal advocacy (EAA), then, is EA in the context of animal advocacy. An effective animal advocate attempts to answer questions like: Can we save or spare more lives if we work on farmed animal welfare or companion animal welfare? Can we save or spare more lives if we take engage in moderate advocacy or radical advocacy? In each case, they attempt to use evidence and reason to answer these questions, and then they pursue the course of action that they estimate will do the most good overall.

This chapter will discuss how effective animal advocates attempt to use evidence and reason to do the most good possible in animal advocacy, as well as conceptual, theoretical, and practical issues that this approach to animal advocacy faces.

2. Conceptual issues

Effective altruists use the following framework to decide which issues to prioritize. First, they think about the *scale* of an issue. If one issue involves more suffering or death than another issue, then it has higher priority for an effective altruist, all else equal. Second, they think about the *neglectedness* of an issue. If one issue is more neglected than another, then it has higher priority for an effective altruist, all else equal. Third, they think about the *tractability* of an issue. If one issue can be addressed more effectively or efficiently than another, then it has higher priority for an effective altruist, all else equal. Finally, they think about *personal fit*. If one issue is a better match for your talents and interests than another, then it has higher priority for you as an effective altruist, all else equal.

What do effective altruists see as the most important issues to be working on right now, considering all these questions? For many effective altruists, the top priority is the category of *existential risk*, i.e., issues that risk

greatly reducing the potential for flourishing in the world (Bostrom 2002). Examples include plagues, pandemics, biotechnology, nanotechnology, nuclear technology, artificial intelligence, and totalitarianism (Bostrom and Cirkovic 2011). These issues are top priority because, in each case, the harm at stake is massive (as well as potentially permanent), neglected, and tractable. Many effective altruists think that we should then prioritize global health and animal welfare, since, in each case, the harm at stake is once again massive (even if not potentially permanent), neglected, and tractable. Of course, many other issues are important too. However, many effective altruists see these as the most important issues to be working on in the present moment.²

What, then, do effective animal advocates see as the most important issues to be working on right now in the context of animal advocacy, considering all these questions? For many effective animal advocates, the top priority is farmed animal welfare. To see why, compare farmed animal welfare with companion animal welfare and wild animal welfare, respectively.³

First, consider how farmed animal welfare compares with companion animal welfare. Farmed animal welfare ranks higher in scale, since many more farmed animals than companion animals suffer and die each year. Farmed animal welfare also ranks higher in neglectedness, since farmed animals receive less support than companion animals do. Finally, farmed animal welfare ranks higher in tractability, since it costs less time, energy, and money to prevent suffering or death for the average farmed animal than it does for the average companion animal. As a result, farmed animal advocacy ranks higher than companion animal advocacy in scale, neglectedness, and tractability. (We will consider objections to this analysis below.)

Now consider how farmed animal welfare compares with wild animal welfare. Wild animal welfare is higher in scale, since wild animals experience even more suffering and death than farmed animals do, as a result of hunger, thirst, illness, injury, predation, and more. Wild animal welfare is also more neglected, since wild animals receive even less support than farmed animals do.⁴ However, wild animal welfare is much *less* tractable, since we currently have no effective and efficient way to help wild animals at all. For many effective animal advocates, the upshot is that we should focus mostly on farmed animal welfare at present, while still focusing on wild animal

² For discussion of these cause areas from an effective altruist perspective, see the Open Philanthropy Project website: <https://www.openphilanthropy.org/focus>

³ For discussion of these cause areas from an effective altruist perspective, see the Animal Charity Evaluators website: <https://animalcharityevaluators.org/advocacy-interventions/prioritizing-causes/#detailed>

⁴ Note that I am not counting environmental conservation work as wild animal welfare work, since this work does not typically address the suffering of individual wild animals.

welfare somewhat. In particular, we should advocate for moral and political standing for wild animals and research interventions in wild animal suffering, so that we can be in a position to help them in the future.⁵ (We will consider objections to this analysis below as well.)

What kinds of approaches do effective animal advocates take to these issues? Consider farmed animal welfare. Effective animal advocates tend to take approaches such as consumer outreach, corporate outreach, and development of alternative food products, such as plant-based meat and cultured meat. In contrast, they do not tend to take approaches such as animal rescue or sanctuary. Why the difference? Many effective animal advocates think that when you engage in consumer outreach, corporate outreach, and development of alternative food products, you are potentially making a difference for relatively many animals at once, and so you are potentially spending less time, energy, and money per animal saved or spared. In contrast, they think that when you engage in animal rescue or sanctuary, you are potentially making a difference for relatively few animals at once, and so you must spend more time, energy, and money per animal saved or spared. (We will consider objections to this analysis below as well.)

Effective animal advocates also aspire to take a pragmatic approach to this work. For example, when it comes to consumer outreach, many effective animal advocates think that if we can save or spare animals more effectively by advocating for reducetarianism (i.e., reduction of animal products) than by advocating for veganism (i.e., elimination of animal products) then we should do so. When it comes to corporate outreach, many effective animal advocates think that if we can save or spare animals more effectively by advocating for, say, cage-free eggs or Meatless Mondays than by advocating for vegan menus, then we should do so. When it comes to development of alternative food products, many effective animal advocates think that if we can save or spare animals more effectively by partnering with conventional food producers than by maintaining full independence, then we should do so. And so on.⁶ (We will once again consider objections to this analysis below.)

Over the past decade or so, a network of organizations has developed around the idea of EA in general and EAA in particular. These organizations are attempting to use evidence and reason to do the most good possible, in part by motivating others to do the same. With respect to EAA in particular, this includes Animal Charity Evaluators (ACE), which is dedicated to finding and promoting the most effective ways to help animals. Among the organizations ACE recommends are Animal Equality, The Good Food Institute, The Humane League, the Albert

⁵ For discussion of wild animal suffering, see Horta 2010, Sebo forthcoming, and Tomasik 2015.

⁶ For arguments in favor of these approaches, see Leenaert 2017. For arguments against, see Taft 2016.

Schweitzer Foundation, Compassion in World Farming, Faunalytics, L214, Otwarte Klatki, ProVeg International, The Nonhuman Rights Project, and Vegan Outreach. Some of these organizations are explicitly committed to EAA. Others are not, but they still do excellent work by EAA standards. I will discuss some of these (and other) organizations below.

This is more or less the interpretation of EAA that many effective animal advocates accept. However, this conception of EAA raises many conceptual questions. Consider two examples.

First, what is the relationship between EAA and different moral theories? Some people see a connection between EAA and utilitarianism, which holds that we morally ought to maximize wellbeing in the world. It makes sense that people would see this connection, since utilitarianism and EAA as I have described it do share several core features. For example, utilitarianism and EAA as I have described it are both committed to *impartial benevolence*, i.e., to helping those in need no matter who they happen to be and no matter how far away they happen to be in space and time. They are both committed to *aggregation*, i.e., to evaluating states of affairs in terms of how much wellbeing they contain. They are both committed to *maximization*, i.e., to producing the best possible state of affairs (though they might or might not agree about what counts as the best possible state of affairs, since they might or might not agree about what constitutes wellbeing). And more. In contrast, other moral theories reject some or all of these features. For example, Kantianism holds that we morally ought to treat others as ends independently of whether doing so produces the best possible outcome. Virtue theory holds that we morally ought to cultivate virtues independently of whether doing so produces the best possible outcome. Care theory holds that we morally ought to cultivate relationships of care independently of whether doing so produces the best possible outcome. And so on.⁷

However, even if EAA as I have described it has more in common with utilitarianism than with other moral theories, a non-utilitarian can take interest in EAA as well. There are at least two reasons for this. First, it is possible to tailor EAA for non-utilitarian theories in much the same kind of way that the standard interpretation tailors it for utilitarianism. After all, all moral theories hold that consequences are morally relevant, even if they hold that other factors are morally relevant too. For example, a Kantian might think that we should, at least in part, promote a world in which others are able to set and pursue ends for themselves. A virtue theorist might think that we should, at least in part, promoting a world in which others can be treated with respect and compassion. A care theorist might think

⁷ For discussion of these moral theories, see Schlottmann and Sebo 2018, Chapter 3.

that we should, at least in part, promoting a world in which others can form and maintain relationships of care. And so on. Moreover, even if these theories do not *require* the pursuit of any particular goal, they might still *permit* the pursuit of particular goals, within certain limits. Granted, these theories might require or permit the pursuit of different goals than utilitarianism, they might require more than the pursuit of these goals, and they might prohibit some means of pursuing these goals. Still, insofar as they require or permit pursuit of particular goals, we can identify a theory-specific interpretation of EAA that endorses informed, rational, effective, efficient pursuit of the goals required or permitted by this theory, within the limits set by this theory.

The second reason that it is possible for a non-utilitarian to take interest in EAA is that it is possible to interpret EAA in more general, theory-neutral terms than I have described here, so that proponents of different moral theories can accept one and the same kind of EAA. For example, on a maximally general interpretation, EAA endorses informed, rational animal advocacy, but does not necessarily endorse any particular goal within animal advocacy, or any particular limit on animal advocacy. This kind of EAA is neutral with respect to all moral theories. However, it risks being too neutral, since it does not take a stand on whether particular goals are good or bad at all. Alternatively, on a moderately general, theory-neutral interpretation, EAA endorses the goal of *minimizing harms*, but does not necessarily endorse the goal of *maximizing benefits* or the idea of minimizing harms *by any means necessary*. This moderately general interpretation strikes a good balance: It allows the EAA framework to take a stand on some issues (e.g., by endorsing the goal of minimizing harms) while leaving other issues open for discussion (e.g., by not taking a stand on the goal of maximizing benefits, or on the need for moral constraints in animal advocacy). This interpretation also coheres with how many people think about triage cases. For example, in a triage case, many people agree that one should attempt to minimize suffering or death in ways that do not violate rights, but people do not necessarily agree about whether one should also attempt to maximize happiness or life, or about whether should attempt to minimize suffering or death in ways that violate rights.⁸

So, which kind of EAA makes the most sense all things considered? This depends in part on how we answer a related conceptual question: What kind of role is EAA meant to play in our thinking about what to do? For example, is EAA meant to be a *theory* that tells us what makes right actions right in the context of animal advocacy? Is it meant to be a *decision procedure that we follow* in our everyday animal advocacy? Is it meant to be a *decision*

⁸ Singer (1972) takes a similar approach by arguing that we have a moral obligation to prevent very bad things from happening if we can do so without sacrificing anything morally significant.

procedure that we promote in our everyday animal advocacy? Is it meant to be a *social movement* that engages in a certain kind of animal advocacy? And so on. Many people talk about EA in general, and about EAA in particular, in a way that combines all of these roles. Yet these roles are importantly different. In particular, if EAA is a theory, then, as with any theory, we should evaluate it in terms of theoretical virtues such as explanatory and predictive power. In contrast, if EAA is a decision procedure that we should follow, promote, or build community around, then, as with any other decision procedure, we should evaluate it in terms of practical virtues such as: Does following, promoting, and/or building community around this decision procedures lead to better outcomes than following, promoting, and/or building community around other decision procedures?

Distinguishing these roles raises an interesting possibility: that we should accept different kinds of EAA for the different roles that it plays in our advocacy. For example, consider how a utilitarian might think about this issue. Insofar as they consider EAA as a theory of what makes right actions right in the context of animal advocacy, they might accept a utilitarian EAA. However, insofar as they consider EAA as a policy that they should follow, promote, or build community around, they will think that it all depends on what does the most good possible. For example, many utilitarians believe that, if our goal is to maximize aggregate wellbeing, then we should not exclusively follow, promote, or build community around utilitarian thinking in our everyday advocacy. Instead, we should follow, promote, and build community around a combination of utilitarian, Kantian, virtue theoretic, and care theoretic terms. Why? Because people who treat others as ends, cultivate virtues, and cultivate relationships of care tend to produce better outcomes than people who do not.⁹

These points are especially important insofar as we think about EAA as a social movement. Here it seems clear that, even if we accept utilitarianism as a moral theory, we should not restrict the EAA movement to fellow utilitarians. Why? Because as Mill (2008) argues, a community that allows members to have different beliefs, values, and practices (within certain limits) tends to function better than a community that requires members to have a single, unified set of beliefs, values, and practices. Granted, a diverse community is more likely to experience conflict and disagreement. But it is also more likely to benefit from an expanded membership, an expanded division of labor, and a healthy sense of epistemic humility. Of course, this does not mean that a utilitarian should accept a

⁹ The idea that we should sometimes follow moral theories other than the one that we accept is called *indirect morality*, and the idea that we should sometimes promote around theories other than the one that we think is correct is called *esoteric morality*. For discussion, see Sidgwick 1981, Hare 1982, Railton 1984, and Parfit 1986.

maximally general interpretation of EAA as a social movement, since it can be good for community members to have at least some shared beliefs, values, and practices. But it might be that a utilitarian should accept a *moderately* general interpretation of EAA as a social movement, for example by basing it on the idea of using evidence and reason to minimize suffering within certain limits, rather than the idea of using evidence and reason to maximize happiness by any means necessary.¹⁰

3. Theoretical issues

I will now summarize some of the main theoretical issues that people have with EAA, i.e. issues that people have with EAA as a theory of what makes right actions right in the context of animal advocacy. I will then summarize the replies that an effective animal advocate can offer to these theoretical critiques. (As we will see, these critiques assume a utilitarian EAA, and I will for the most part do the same in my responses.)

Here, then, are four common theoretical issues that people have with EAA.

1. *We should not evaluate states of affairs in terms of aggregate wellbeing.* First, some critics take issue with the idea that we should evaluate states of affairs in terms of how much wellbeing they contain. In particular, if we evaluate states of affairs this way, then we will sometimes produce surprising results. For example, we might find that reducing farmed animal suffering or wild animal suffering by 1% is better than reducing companion animal suffering by 99%. After all, there are so many farmed animals and wild animals in the world that a minor reduction in suffering per farmed animal or wild animal outweighs a major reduction in suffering per companion animal in the aggregate. Similarly, we might find that reducing insect suffering by 1% is better than reducing human suffering by 99%. After all, there are about one billion insects for every human in the world. Thus, if we suppose that the average insect experiences even one millionth the amount of wellbeing that the average human does at any given time (a conservative assumption in my view), it follows that the ratio of insect wellbeing to human wellbeing in the world at any given time is still about one thousand to one, in which case a 1% reduction of suffering per insect does, indeed, outweigh a 99% reduction of suffering per human in the aggregate. Yet one might find this kind of aggregation

¹⁰ Whether proponents of other moral theories can think about this issue in the same kind of way depends on whether they allow for indirect or esoteric morality. If they do not, then they would be committed to following and promoting the principle they accept as theoretically correct, independently of whether doing so is practically best.

implausible. In particular, one might think that this kind of aggregation ignores the inherent moral value of individual human and nonhuman animals, as well as the inherent moral value of fairness, equality, and other such distributions of wellbeing in the world.¹¹

2. *We should not harm the few for the sake of the many.* Second, some critics take issue with the idea that we should harm the few so that we can help the many. For example, suppose that we face a choice between promoting reductarianism and promoting veganism. If we stipulate that promoting reductarianism sometimes saves or spares more lives than promoting veganism does, then utilitarian EAA implies that we are sometimes morally required to promote reductarianism, in spite of the fact that doing so risks appearing to support the idea that animals are food.¹² Similarly, suppose that we face a choice between using racism, sexism, and ableism in animal advocacy (for example, by making simplistic, reductive, and appropriative comparisons between human and nonhuman oppressions) and refusing to do so (for example, by refusing to make such comparisons). If we stipulate that the former approach sometimes saves or spares more lives than the latter approach, then utilitarian EAA implies that we are sometimes morally required to take the former approach, in spite of the fact that doing so risks appearing to support the idea that some humans are “less-than-human”. Yet one might find these results implausible. In particular, one might think that we are morally required not to risk supporting the idea that some animals are food or that some humans are “less-than-human” independently of whether these approaches do the most good possible.¹³ (Note that we will later be questioning these assumptions about what does the most good possible.)

3. *We do not need to neglect the few for the sake of the many.* Third, some critics take issue with the idea that we should neglect helping the few so that we can help the many instead. For example, suppose that we face a choice between supporting rescue or sanctuary, on one hand, and supporting farmed or wild animal advocacy, on the other hand. If we stipulate that supporting farmed or wild animal advocacy saves or spares more lives than supporting rescue or sanctuary does (at least sometimes), then utilitarian EAA implies that we are morally forbidden from supporting rescue or sanctuary (at least sometimes), in spite of the fact that this activity would help many animals.¹⁴ Similarly, and as above, suppose that we face a choice between multi-issue animal advocacy (i.e., animal

¹¹ For discussion of the ethics of aggregation, see Voorhoeve 2014 and Temkin 1996. For discussion of insect suffering, see Horta 2010, Sebo forthcoming, and Tomasik 2015.

¹² For discussion of reductarianism, see Kateman 2017.

¹³ For discussion of racism in animal advocacy, see Harper 2010. For discussion of sexism in animal advocacy, see Gaarder 2011. For discussion of ableism in animal advocacy, see Taylor 2017.

¹⁴ For discussion of the moral and political value of sanctuaries, see Donaldson and Kymlicka 2015.

advocacy that aspires to stand in solidarity with other social movements) and single-issue animal advocacy (i.e., animal advocacy that does not aspire to stand in solidarity with other social movements). If we stipulate that a single-issue approach sometimes saves or spares more lives than a multi-issue approach does, then utilitarian EAA implies that we are sometimes morally forbidden from taking a multi-issue approach, in spite of the fact that this activity would help many humans. Yet one might find these results implausible too. In particular, one might think that we are morally permitted to support rescue and sanctuary and other social movements independently of whether this approach does the most good possible.¹⁵ (Note that, as above, we will later be questioning these assumptions about what does the most good possible.)

4. *We do not need to demand so much from ourselves.* Fourth, and relatedly, some critics take issue with EAA because they see it as too demanding, physically as well as psychologically. First, they see EAA as too *physically* demanding because, as we have seen, if our aim is to do the most good possible, then we need to apply this framework to every aspect of our lives. We need to apply it to what major we select in college, what job we accept after college, how we spend our nights and weekends, and so on. Moreover, while we might hope that living a relatively happy life will allow us to produce better consequences than living a relatively unhappy life will, ultimately, we are morally required to do produce the best possible consequences independently of how happy we are, according to utilitarian EAA. Second, critics see EAA as too *psychologically* demanding because, as we have seen, if our aim is to do the most good possible, then the rightness of every action we perform depends on the product of a complicated calculation. We have to ask: Which action, of every action available to me, will maximize aggregate wellbeing from now until the end of time? This question is difficult if not impossible to answer with confidence. Thus, it is difficult if not impossible to know which of our actions, if any at all, are morally right and wrong according to utilitarian EAA.¹⁶ Yet one might find these results implausible too. In particular, one might think that there is a limit to how demanding morality can be. We should be able to be relatively happy in life, and we should be able to know which of our actions, if any at all, are morally right and wrong, within certain limits.¹⁷

With respect to all four of these theoretical objections, an effective animal advocate has at least three replies available, which are all compatible with each other. The first reply, which relates to our conceptual

¹⁵ For discussion of the ethics of multi-issue advocacy, see Sebo 2018.

¹⁶ For discussion, see Kagan 1991, Lenman 2005, Railton 1984, and Scheffler 1982.

¹⁷ For discussion of what life is like for people who try to live altruistically, see MacFarquhar 2016.

discussion above, is to reject a utilitarian EAA, and, therefore, to deny that EAA faces these objections. In particular, if one accepts, say, a Kantian, virtue theoretic, or care theoretic EAA, then this kind of EAA might or might not face these objections. For example, a Kantian EAA does not (1) require us to evaluate states of affairs in terms of aggregate wellbeing, (2) require us to harm one for the sake of helping five, (3) require us to neglect one for the sake of helping five, or (4) demand that we spend all our time thinking about what will do the most good, and then acting accordingly. It simply instructs us to use evidence and reason to effectively and efficiently pursue morally permissible ends within animal advocacy. Likewise, if one accepts a general, theory-neutral EAA, then this kind of EAA might or might not face these objections either, for the same reason. For example, an EAA that does not take a stand on what goal to pursue or whether to place moral limits on our pursuit of this goal will not face these objections to the same degree, since it will not, in and of itself, require as many seemingly wrong actions.

However, we can note two caveats about this kind of reply. First, this reply will not work for people who accept a utilitarian EAA. Second, it will not allow an effective animal advocate to avoid theoretical issues. All it does is replace some theoretical issues with others, or make some theoretical issues “external” to EAA rather than “internal” to EAA. In particular, if one accepts, say, a Kantian, virtue theoretic, or care theoretic EAA, all this does is replace the kinds of theoretical issues that utilitarianism faces with the kinds of theoretical issues that these other theories face. For example, a Kantian EAA will seem to be too restrictive in some cases (since it forbids harming others even when we think we should) and too permissive in others (since it permits not helping others even when we think we should). Second, and relatedly, if one accepts a general, theory-neutral EAA, all this does is make some theoretical issues “external” to EAA rather than “internal” to EAA. For example, one will still have to decide whether to maximize happiness or accept moral constraints in animal advocacy. The only difference is that these decisions will now be about how to apply EAA than about how to interpret EAA. Thus, no matter what kind of EAA we select, we will not be able to avoid theoretical issues, and so the task is not to try to do so. We will return to this point in a moment.

The second reply available to an effective animal advocate is to assert that, even if we accept a utilitarian EAA, this kind of EAA does not have these implications as often as we might fear. In particular, as we discussed above and will discuss in more detail below, we might think that if our goal is to do the most good possible in animal advocacy, then we should often avoid (1) evaluating states of affairs in terms of aggregate wellbeing, (2) harming the few for the sake of helping the many, (3) neglecting the few for the sake of helping the many, and (4)

spending all our time thinking about how to do the most good possible and then acting accordingly. Why? Because while these activities might be optimal in the short term, they are arguably suboptimal in the long run, since they prevent us from following the kinds of rules, cultivating the kinds of virtues, and cultivating the kinds of relationships of care that arguably do the most good possible in the long run. (We will consider these ideas further in the next section.)

However, we can note two caveats about this kind of reply as well. First, even if thinking like a utilitarian is often suboptimal, it is not necessarily always suboptimal. For example, it might at least *sometimes* be the case that violating individual rights does the most good possible in practice, or that sacrificing our own wellbeing does the most good possible in practice. Insofar as it is, utilitarian EAA implies that we should pursue these courses of action in practice. Second, no matter how often utilitarian EAA has these implications in practice, the fact remains that it has these implications in theory. That is, a utilitarian EAA is committed to the idea that *if* violating individual rights does the most good possible, or *if* sacrificing our own wellbeing does the most good possible, *then* we should pursue these courses of action. For many people, that will be reason enough to reject utilitarian EAA in favor of a different interpretation of EAA. (We will consider these caveats further in the next section as well.)

The third reply available to an effective animal advocate is to accept that, insofar as utilitarian EAA has these implications in practice, these implications are correct. Granted, it might initially seem implausible that we should, say, violate individual rights or sacrifice our own wellbeing for the greater good. But recall that every moral theory has implausible implications. The point of practical ethics is to consider the pros and cons of every moral theory holistically, so that we can select the moral theory that makes the most sense overall. Moreover, at least some implications that seem implausible at first might seem more plausible on reflection. For example, it might initially seem implausible that we should sometimes violate individual rights, but this idea will start to seem more plausible if we start to think of everyday life as an emergency situation that requires tragic choices rather than as a mundane situation that does not. Similarly, it might initially seem implausible that we should sometimes sacrifice our own wellbeing, but this idea might start to seem more plausible if we start to think of everyday life as an emergency situation that requires self-sacrifice, rather than as a mundane situation that permits full individual liberty.

My view is that these three replies, suitably fleshed out, are compelling. They are especially compelling when we think about them holistically. Insofar as we accept a non-utilitarian EAA, we can avoid these objections (but will then face others). And, insofar as we accept a utilitarian EAA, we can reply to these objections in the same

kind of way that a utilitarian can, i.e. by arguing that EAA does not always have these implications in practice, and that when it does, these implications are correct. (Likewise, if we accept a non-utilitarian EAA, then we can reply to the theoretical objections that it faces in the same kind of way, by arguing that EAA does not always have these implications in practice, and that when it does, these implications are correct.) Granted, utilitarian EAA has revisionary and demanding implications for animal advocacy. Then again, why should we have expected otherwise? The world contains vast amounts of suffering that we can reduce substantially, if only we are willing to commit to the task. And, if committing to the task means sometimes deviating from what our often uninformed, irrational, and outdated moral intuitions permit or require, then so be it.

4. Practical issues

I will now summarize some of the main practical issues that people have with EAA, i.e. issues that people have with EAA as a decision procedure to follow, promote, or build community around in animal advocacy. I will then summarize the replies that an effective animal advocate can offer to these practical critiques. (As we will see, these critiques once again assume a utilitarian EAA, and I will for the most part do the same in my responses.)

The practical concerns that we will consider here are all based on the idea, considered briefly above, that EAA thinking is not always the best way to achieve EAA aims. On this view, if we want to do the most good possible in animal advocacy, then we should not spend all of our time, energy, and money pursuing this goal directly. Instead, we should spend some time, energy, and money pursuing this goal directly and some time, energy, and money pursuing it indirectly, by pursuing other goals directly instead. Only then will we be able to create the kind of individual and collective structures that will allow us to maximize aggregate wellbeing in the world.

In particular, consider four respects in which one might worry that EAA thinking is not always the best way to achieve EAA aims (which mirror the theoretical issues that we discussed in the previous section).

1. Evaluating states of affairs in terms of aggregate wellbeing is not always optimal. First, EAA thinking invites a focus on aggregate wellbeing rather than on rights or justice. Yet, one might worry, this focus on aggregate wellbeing can be self-defeating in some cases in practice, since it might be that promoting rights or justice in the short term is necessary for promoting aggregate wellbeing in the long run. This could be true for many reasons. For example, it might be that if people see animals as mere “receptacles of utility” rather than as individuals who merit

rights and justice, then people will be less likely to promote particular kinds of social, legal, and political changes for animals (e.g., they will be less likely to support moral personhood, legal personhood, and political membership for animals). And it might be that if we want to maximize aggregate wellbeing for the long run, then we should promote these kinds of social, legal, and political changes in the short term. This kind of reasoning, if correct, would support placing moral limits on aggregation in some contexts.¹⁸

2. *Harming the few for the sake of the many is not always optimal.* Second, EAA thinking invites a willingness to harm the few for the sake of helping the many all else equal. Yet, one might worry, this kind of moral math can be self-defeating in some cases in practice, since it might be that refusing to harm the few in the short term is a necessary means to helping the many in the long run. This might be true in many contexts. For example, it might be that if we often harm some animals to help other animals (e.g., by promoting reducetarianism), then we will undermine efforts to promote rights or justice for animals, which might be necessary for maximizing aggregate wellbeing in the long run. Similarly, it might be that if we often harm humans to help nonhumans (e.g., by promoting animal rights in racist, sexist, or ableist ways), then we will undermine efforts to promote rights and justice for all (as well as to build coalitions across movements), which might also be necessary for maximizing wellbeing in the long run. This kind of reasoning, if correct, would support placing moral limits on animal advocacy in some contexts.

3. *Neglecting the few for the sake of the many is not always optimal.* Third, EAA thinking invites a willingness to neglect the few for the sake of helping the many all else equal. Yet, one might worry, this kind of moral math can be self-defeating in some cases in practice, since it might be that helping the few in the short term is a necessary means to helping the many in the long run. This might be true in many contexts as well. For example, it might be that if we never provide support for individual animals (e.g., if we never support rescue or sanctuary), then we will undermine efforts to show humans a different, more equitable kind of relationship with nonhumans, which might be necessary for maximizing aggregate wellbeing in the long run. Similarly, it might be that if we never provide support for humans (e.g., if we never stand in solidarity with other social movements), then, as before, we will undermine efforts to promote rights and justice for all (as well as to build coalitions across movements), which

¹⁸ For discussion of a related conception of animal dignity, see Gruen 2014.

might be necessary for maximizing wellbeing in the long run. This kind of reasoning, if correct, would support embracing a plurality of approaches within animal advocacy in some contexts.

4. *Demanding so much from ourselves is not always optimal.* Finally, EAA thinking invites a willingness to neglect our own wellbeing for the sake of the aggregate wellbeing. Yet, one might worry, this kind of approach to animal advocacy might be self-defeating in some cases in practice, since it might be that creating at least some space in life for other pursuits is necessary for doing the most good possible in the long run. After all, we all have limited time, energy, and information, and so we will not always be able to make choices by calculating the expected consequences of all possible courses of action, and then selecting the course of action whose expected consequences are best. Moreover, if we spend all our time, energy, and money attempting to do the most good possible, then we will likely not invest in the kinds of personal and professional projects and relationships that will make advocacy sustainable for us in the long run. This kind of reasoning, if correct, would support relaxing certain kinds of demands on animal advocacy (and life more generally). It would also support investing in self-care and community care as part of EAA, in spite of the fact that the benefits of such work can be difficult to measure.¹⁹

As discussed in the previous section, even if we accept these points in some cases, we might not accept them in all cases. That is, we might think that there are many cases where EAA thinking is optimal (and so we should attempt to do the most good possible directly) and many cases where EAA thinking is suboptimal (and so we should attempt to do the most good possible indirectly, by pursuing other goals directly instead). If so, then the challenge is to learn how to strike a virtuous balance in practice. We will return to this point in a moment.

Why might one worry that EAA thinking will produce these kinds of suboptimal outcomes? There are at least two related reasons. First, there is arguably a risk of methodological bias in EAA thinking. In particular, if your aim is to use evidence and reason to do the most good possible, there is a risk that you will engage in explicit cost-benefit analysis more often than you should (i.e. that you will use this decision procedure when other decision procedures would be better). And, insofar as you use explicit cost-benefit analysis more often than you should, there is also a risk that you will develop a bias in favor of interventions whose benefits are relatively easy to measure or estimate and/or whose costs are relatively hard to measure or estimate, and against interventions whose benefits are relatively hard to measure or estimate and/or whose costs are relatively easy to measure or estimate.

¹⁹ For discussion of self-care and community care in activism, see Jones 2007.

Second, there is arguably a risk of standpoint bias in the current EAA movement. In particular, the current EAA movement has a relatively high percentage of utilitarians and other consequentialists. It also exists at the intersection of the effective altruism and animal advocacy movements, both of which have a reputation (partly earned, partly not) of being for wealthy white people. Moreover, the EAA movement functions as a social community too. This has many benefits, since altruism is easier when done in community with others. However, it also has risks, since it can create a social barrier between the EAA movement / community and other movements / communities. This, in turn, can make it hard for the EAA movement / community to expand its membership, as well as for it to collaborate and communicate effectively with other movements / communities.

The worry, then, is that these methodological and demographic issues might conspire to create a risk of bias in favor of, for example, working within current systems to bring about moderate reforms. And, at least to a degree, the EAA movement does seem to have this priority. However, we need to be careful here. First, even if this priority makes sense in light of bias, that does not mean that it is a product of bias. Second, even if it is a product of bias, that does not mean that it is suboptimal. Third, other movements and communities are susceptible to bias as well, so it is not as though risk of bias is unique to EAA. Fourth, one virtue of EAA is a commitment to taking seriously the effects that biases and heuristics can have on our thinking about what to do. So insofar as effective animal advocates detect a risk of bias, they will be motivated to try to mitigate that risk.

Supposing, then, that the risk of methodological bias and standpoint bias are real, what can effective animal advocates do to preserve the benefits of this framework while mitigating the risks? Here are two complementary steps that they can take, both of which many are already taking.²⁰

First, effective animal advocates can work to *improve* their use of EAA thinking. They can strike a balance between quantitative and qualitative assessments, as well as between assessments of individual interventions and of sets of interventions. For example, they can study the history of social movements to see what sets of interventions tend to work well and what sets of interventions tend to work badly. Then, with respect to the sets of interventions that seem to be effective for other movements. They can then study the degree to which the animal protection movement is relatively similar, and the degree to which the animal protection movement is taking a relatively similar approach. Of course, this will not eliminate the risk of bias. But insofar as effective animal advocates strike a

²⁰ The remainder of this section draws from, and overlaps with, the suggestions made in Sebo and Singer 2018.

virtuous balance between quantitative and qualitative reasoning and between individualistic and holistic reasoning, it will at least reduce this risk of bias.

Second, and relatedly, effective animal advocates can work to *supplement* their use of EAA thinking. They can strike a balance between cost-benefit analysis and other decision procedures in everyday life. In particular, insofar as we have reason to think that engaging in cost-benefit analysis will be useful, we should use this decision procedure. Insofar as we do not, we should use other decision procedures. For example, Kantian thinking can help us to establish valuable rules to follow. Virtue theoretic thinking can help us to cultivate valuable character traits. Care theoretic thinking can help us to cultivate valuable relationships within and outside the movement. Effective animal advocates may or may not think that these moral theories are correct. But even if they do not, they can still benefit epistemically and practically from using these theories as practical guides in certain cases.

As part of these efforts, effective animal advocates work to promote diversity, equity, and inclusion within the movement and links across movements. This can take the form of pursuing unity (working with people in other movements), solidarity (helping / not harming people in other movements), and mutual understanding (sharing information and arguments with people in other movements). This will allow effective animal advocates to expand the demographics of the movement and build coalitions with other movements – all of which will benefit the EAA movement epistemically (by allowing it to incorporate a wider range of perspectives) and practically (by allowing it to incorporate a wider range of skills).²¹ This, in turn, can help mitigate the risk of methodological bias and standpoint bias, since other methodologies and standpoints will start to become more prevalent within the movement.

Fortunately, many effective animal advocates are working to improve and supplement their application of direct utilitarian reasoning in exactly these ways. For example, Animal Charity Evaluators, Faunalytics, Sentience Institute, and other such nonprofit organizations are conducting and promoting quantitative as well as qualitative research about particular interventions as well as sets of interventions. This includes research about farmed animal advocacy and wild animal advocacy, confrontation and conciliation, institutional change and individual change, and more. Moreover, organizations such as Encompass are working to promote diversity, equity, and inclusion within

²¹ For discussion of the epistemic benefits of diversity, see Kitcher 1990 and Strevens 2003.

the EAA movement, such that it can retain its focus on impartially benevolent farmed animal advocacy while expanding and diversifying the range of people who hold positions of power within the movement.

Of course, there are risks involved with some of this work. For example, insofar as effective animal advocates promote non-utilitarian decision procedures and promote viewpoint diversity within the movement, they will likely invite disagreement into the movement about what to do and how to do it. People will disagree more about facts (e.g. what kind of political system would do the most good) as well as about values (e.g. whether we should be trying to do the most good). And while disagreement can have benefits (e.g., it can result in epistemic progress for all involved), it can also have risks. For example, it can result in conflict. It can result in value drift, i.e., in the EAA movement deviating from its current mission. It can also result in fragmentation, i.e., in the EAA movement fragmenting into multiple movements with multiple missions. It can even result in dissolution, i.e., in the EAA movement dissolving entirely.

How should we assess these risks and benefits? Especially given the possibility of bias, we should proceed carefully. On one hand, conflict, value drift, fragmentation, or dissolution would clearly have bad effects. On the other hand, they might not happen, and even if they did happen, they could also have good effects. Conflict could lead to necessary epistemic and practical change. Value drift could lead to a mission that preserves the current strengths of EAA while overcoming the current limitations. Fragmentation could lead to multiple movements that accomplish more collectively than a single movement could accomplish individually. Dissolution could lead current effective animal advocates to do similar work in the context of animal advocacy in general, whether or not this work is still labeled as EAA.

I will not try to reach an overall assessment of these risks and benefits here. Instead, I will simply express my own general hope for the future of EAA. In my view, the ideal EAA movement would resemble the current EAA movement in many ways, e.g., it would accept utilitarian ideals and use cost-benefit analysis to pursue these ideals whenever possible. However, it would also improve and supplement cost-benefit analysis in the ways discussed above, so that the EAA movement can expand and diversify, can communicate and collaborate more effectively with other movements, and can allow for creativity, improvisation, and toleration of difference, within certain limits.

If this is right, the upshot is that the EAA movement should strive to be (a) *much* more utilitarian than the current animal advocacy movement is and (b) at least a *bit* less utilitarian than the current EA movement is.

However, as discussed above, this is likely to be a contextual matter. For example, it might be that EAAs should

follow a more utilitarian decision procedure than they should promote or build this movement around (though even the latter should likely preserve core features of utilitarianism). Similarly, it might be that effective animal advocates should build the movement around a more utilitarian decision procedure at present, given that EAA represents a relatively small subset of the animal protection movement at present (and so can complement non-utilitarian work that others are doing), than in the future, when EAA might represent a larger subset of the animal protection movement, and so will need to create space for complementary work internally (though here too, I think it will be important to preserve core features of utilitarianism).

5. Conclusion

It is tempting to simply think of EAA as utilitarian animal advocacy, and to accept or reject it accordingly. But the truth is more complex. As with all social movements, the meaning of the EAA movement is dynamic and contested. Even at the theoretical level, it is not clearly restricted to utilitarianism, since many effective animal advocates reject aspects of utilitarianism. And, especially at the practical level, it is clearly not restricted to utilitarianism, since effective animal advocates might need to think in non-utilitarian terms in many cases in order to promote utilitarian aims. Nevertheless, I think it makes sense to see utilitarianism as an ideological center of gravity for EAA, since they share not only share a history but also a commitment to informed, rational, impartial benevolence for animals. Insofar as they share this commitment, I think that they are right to do so, and I hope that utilitarians and non-utilitarians alike will work together to pursue this ideal.

Acknowledgements

Thanks to Bob Fischer and Tyler John for helpful feedback and discussion.

References

Bostrom, Nick, 2002. "Existential risks," *Journal of Evolution and Technology* 9:1.

Bostrom, Nick and Milan Cikovic, eds., 2011, *Global Catastrophic Risks* (Oxford: Oxford University Press).

Donaldson, Sue and Will Kymlicka, 2015. "Farmed animal sanctuaries: the heart of the movement?" *Politics and Animals* 1: 50-74.

Gaarder, Emily. 2011. *Women and the Animal Rights Movement* (New Brunswick: Rutgers University Press).

Gruen, Lori. 2014. "Dignity, captivity, and an ethics of sight." in Lori Gruen, ed., *The Ethics of Captivity* (Oxford: Oxford University Press): 231-247.

Hare, R. M. 1982. *Moral thinking: its levels, methods and point* (Oxford: Oxford University Press).

Horta, Oscar, 2010. "Debunking the idyllic view of natural processes: population dynamics and suffering in the wild," *Télos* 17, 73-88.

Harper, Breeze (ed.), 2010, *Sistah Vegan: Black Female Vegans Speak on Food, Identity, Health, and Society* (New York: Lantern Books).

jones, patrice. 2007. *Aftershock: confronting trauma in a violent world* (New York: Lantern Books).

Kagan, Shelly. 1991, *The Limits of Morality* (Oxford: Clarendon Press).

Kateman, Brian (ed.), 2017, *The Reductarian Solution* (New York: TarcherPerigee).

Kitcher, Philip, 1990, "The division of cognitive labor," *Journal of Philosophy* 87: 5-22.

Leenaert, Tobias, 2017, *How to create a vegan world: A pragmatic approach* (New York: Lantern Books).

Lenman, James. 2005. "Consequentialism and Cluelessness," *Philosophy & Public Affairs* 29:4: 342-370.

MacAskill, William, 2015, *Doing good better: How effective altruism can help you make a difference* (New York: Gotham Books).

MacFarquhar, Larissa. 2016. *Strangers drowning: impossible idealism, drastic choices, and the urge to help* (New York: Penguin).

Midgley, Mary. 1998. *Animals and why they matter* (Athens: The University of Georgia Press).

Mill, John Stuart. 2008. *On liberty and other writings* (Cambridge: Cambridge University Press).

----- 2004. *Utilitarianism and other essays* (London: Penguin).

Parfit, Derek, 1986, *Reasons and Persons* (Oxford: Oxford University Press).

Railton, Peter, 1984, "Alienation, consequentialism, and the demands of morality," *Philosophy and Public Affairs* 13:2: 134-71.

Reese, Jacy, 2018, *The End of Animal Farming* (Boston: Beacon Press).

Samuel Scheffler, 1982, *The Rejection of Consequentialism* (Oxford: Oxford University Press).

Schlottmann, Christopher and Jeff Sebo, 2018, *Food, Animals, and the Environment: An ethical approach* (New York: Routledge).

Sebo, Jeff and L.A. Paul, forthcoming, "Effective Altruism and Transformative Experience," in Hilary Greaves and Theron Pummer, eds., *Philosophy and Effective Altruism* (Oxford: Oxford University Press).

Sebo, Jeff and Peter Singer, 2018, "Activism," in Lori Gruen, ed., *Critical Terms for Animal Studies* (Chicago: Chicago University Press).

Sebo, Jeff, 2018, "Multi-issue food activism," in Anne Barnhill, Tyler Doggett, and Mark Budolfson, eds., *The Oxford Handbook of Food Ethics* (Oxford: Oxford University Press).

Sebo, Jeff, forthcoming, "Animals and climate change," in Mark Budolfson, Tristram McPherson, and David Plunkett (eds.), *Philosophy and Climate Change* (Oxford: Oxford University Press).

Sentience Institute, 2018, "Summary of Evidence for Foundational Questions in Effective Animal Advocacy": <https://www.sentienceinstitute.org/foundational-questions-summaries>

Sidgwick, Henry, 1981, *The Methods of Ethics* (Indianapolis: Hackett).

Singer, Peter. 1972. "Famine, affluence, and morality." *Philosophy and Public Affairs* 1:1: 229-243

Singer, Peter, 2015. *The most good you can do: how effective altruism is changing ideas about living ethically* (Yale: Yale University Press).

Strevens, Michael, 2003. "The role of the priority rule in science," *Journal of Philosophy*, 100:2: 55-79.

Taft, Casey, 2016, *Motivational methods for vegan advocacy: A clinical psychology perspective* (Danvers: Vegan Publishers).

Taylor, Sunaura. 2017. *Beasts of Burden: Animal and Disability Liberation* (New York: New Press).

Temkin, Larry, 1996. "A continuum argument for intransitivity." *Philosophy & Public Affairs* 25:3: 175-210.

Tomasik, Brian, 2015. "The importance of wild animal suffering," *Relations* 3:2: 133-52.

Vinding, Magnus, 2018, *Effective altruism: How can we best help others?* (The Neuroethics Foundation).

Voorhoeve, Alex, 2014, "How should we aggregate competing claims?" *Ethics* 125:1: 64-87.