

# EFFECTIVENESS AND ECUMENICITY

Chong-Ming Lim

## Abstract

Effective altruism is purportedly ecumenical towards different moral views, charity causes, and evidentiary methods. I argue that effective altruists' criticisms of purportedly less ineffective charities are inconsistent with their commitment to ecumenicity. Individuals may justifiably support charities other than those recommended by effective altruism. If effective altruists take their commitment to ecumenicity seriously, they will have to revise their criticisms of many of these charities.

## 1 Introduction

Effective altruism is a growing philosophical and philanthropic movement. Its main draw is its claim to be 'based on a very simple idea: we should do the most good we can' (Singer 2015: vii) by using careful reasoning and good evidence. Based on this idea, effective altruists have produced impressive evaluations and rankings of many existing charities, in terms of how well (or effectively) they save and improve the lives of others (especially those in relatively poorer areas of the world). Individuals – especially those in relatively well-off (global) positions – are urged to donate to the most effective charities. Some effective altruists have also recommended individuals to choose their careers based on what would generate the highest income, so as to best fulfil the demand to do the most good (MacAskill 2014). Effective altruists have done much good. Many lives have been saved or improved by donations inspired by effective altruism. They have also prompted a more general reckoning of the aims and means of philanthropy, and rejuvenated discussions of what constitutes an ethical person or life.

Yet since its earliest articulation, the implications of this simple idea have been severely criticised. Effective altruism is accused of focusing on the symptoms of inequality and suffering, to the neglect of political or institutional solutions; of attending to short- rather than long-term effects; of perpetuating the existence of the very structures that generate the problems they wish to resolve; of ignoring or requiring us to sacrifice other moral considerations beyond saving and improving lives; among others (Herzog 2016; Rubenstein 2016; Gabriel 2017). More generally, many people are also put off by what they regard as the arrogance of some effective altruists, as conveyed by their sweeping criticisms of many individuals and charities.

Partly in response to these criticisms, proponents of effective altruism have made three ecumenical concessions. First, individuals may endorse a whole range of different moral views beyond the consequentialism or welfarism which have been described as undergirding effective altruism. Second, and beginning from different moral views, individuals may endorse a whole range of different causes and charities which promote different values. Third, individuals and charities may employ different methods to assess whether their activities are effective. These are sensible responses, and go some way in accounting for the complexities of philanthropy. They have, however, received relatively little philosophical attention. I argue

that effective altruists' criticisms of purportedly less ineffective charities are inconsistent with their commitment to ecumenicity. Individuals who choose to support these charities are not misguided. Instead, they may simply be disagreeing with effective altruists' starting moral views and endorsed methods. Effective altruists' criticisms will have to be revised.

The paper unfolds as follows. In Section 2, I distinguish between thick and thin effective altruism, and situate the three ecumenical moves in the context of criticisms that thick effective altruism has received. In Section 3, I employ these moves to re-evaluate effective altruists' criticisms of certain charities. My aim is not so much to deliver a substantive objection against effective altruism, as it is to consider the implications of the three ecumenical concessions for the criticisms made by effective altruists of some purportedly less effective charities. In Section 4, I briefly consider potential objections by effective altruists, which begin from the distinction between thick and thin effective altruism. I conclude in Section 5.

## 2 Commitment and Ecumenicity

Iason Gabriel distinguishes between two versions of effective altruism. On the thin version, effective altruism holds simply that individuals should do the most good, on the basis of good evidence. In practical terms, this directs individuals to donate a significant amount of their resources to the most effective charities. This claim is compatible with a 'wide range of moral theories and remains noncommittal both about the nature of the good and about the individual's relationship to it' (Gabriel 2017: 458). On the thick version, however, effective altruism 'makes a number of further assumptions' (2017: 458). Instead of being noncommittal among different moral views, thick effective altruism adopts a broadly consequentialist moral theory. It also adopts a welfarist understanding of value, according to which states of affairs are evaluated and ranked by the amount of suffering and deaths they contain. It also uses specific scientific tools, such as 'cost-effectiveness analyses and randomisation [tests] to help quantify and compare the impact of different interventions' (2017: 459). While Gabriel recognises that not all effective altruists endorse thick effective altruism, he nonetheless centres his discussion on the latter. For him, the commitments of thick effective altruism 'explain many of its judgments and capture much of what makes it unique' (2017: 459).

As Gabriel notes, many criticisms beset thick effective altruism. I will not rehearse all of these, but will briefly present only those which help to situate my subsequent discussions.

Thick effective altruism's commitment to welfarism is taken to imply that it is not intrinsically good when values such as equality or justice are realised; they are good only insofar as they have a positive effect on welfare. For many people, this runs counter to their considered judgements – often arising from holding moral views which are not welfarist, or not entirely welfarist – that values such as equality do have intrinsic value (2017: 459). More generally, the worry is that thick effective altruism requires individuals to forgo those of their concerns and commitments which are not fully accounted for by welfarism. For instance, Peter Singer counsels individuals who are passionate about the arts to instead direct their resources to other, more effective charities that save or improve more lives (Singer 2015: 118-120). This commitment is criticised for neglecting the fact that there is more of what is morally valuable within human and social life, than that which is identified by the moral view which effective altruism adopts (Gray 2015; Krishna 2016).

Gabriel also criticises this commitment, along with their adopting a broadly consequentialist moral view, for leading effective altruists to focus only on causes that promote welfarist values (construed a certain way), to the neglect of other causes. Combined with effective altruists' commitment to doing the *most* good, this leads to broader causes such as institutional reform being side-lined. This is because institutional reform is more complicated and difficult, takes a much longer time to achieve, and often involve greater costs than direct interventions to save or improve lives. Yet unjust institutional structures are exactly the root of the suffering that effective altruists try to address. In effect, effective altruists' concern with the latter is short-sighted, addressing the symptoms rather than the cause (Clough 2015; Srinivasan 2015).<sup>1</sup>

Effective altruists have also been criticised for what appears to be an over-reliance on a specific method of collecting and assessing evidence – such as through Randomised Controlled Trials (RCTs).<sup>2</sup> In RCTs, certain features are held constant while others are varied – as in a controlled experiment. For instance, effective altruists may measure the effect of distributing insecticide-treated bed-nets by comparing places which receive such nets, to those which do not. This allows them to clearly distinguish the effect of any specific direct intervention, from those due to other sources. However, critics argue that RCTs are limited in their scope of application. While they are appropriate for assessing interventions for which many variables may be easily manipulated. They provide very clear evidence only within a short time-frame. RCTs are also 'bad at detecting any unintended effects of a programme, especially those effects that fall outside the population or timeframe that the organization or researchers had in mind' (Clough 2015: 2-3). In effect, RCTs are not well-suited for evaluating large-scale projects that span across a long time – such as institutional reform. Effective altruists' reliance on RCTs means that they may appropriately evaluate the effectiveness of only small-scale and short-term interventions.

Given these criticisms, Gabriel argues that effective altruists have to be more accommodating of different moral views and values beyond those which thick effective altruism is committed to, and of different standards of assessing evidence (Gabriel 2017: 470). In response to Gabriel, several effective altruists have drawn a distinction between the 'definition of effective altruism', and the 'actions and recommendations of the effective altruist community' (Halstead et al 2017: 3-4). They agree that some effective altruists (and the organisations to which they belong) are indeed committed to narrower moral and methodological views. These thick effective altruists may indeed be criticised along the lines Gabriel sketches. However, they claim that arguments against these thick effective altruists are 'not necessarily valid against the idea of effective altruism as expressed by its definition' (2017: 3-4). The definition of effective altruism is simply that individuals must use 'evidence and reason to figure out how to benefit others as much as possible, and taking action on that basis' (2017: 4). This is thin effective altruism, and 'allows for a wide range of ethical aims and methods to achieve those aims' (2017: 4). Thin effective altruism 'is ecumenical between a range of moral theories (2017: 7), and neutral among different causes and methodological tools (2017: 11-12). Individuals and charities may hold different (and even competing) moral views, or be committed to different (and even competing) causes, *yet* be considered as part of

---

<sup>1</sup> As Jeff McMahan (2017) notes, similar criticisms were made much earlier by Bernard Williams (1973) and Martha Nussbaum (1997), against utilitarianism and its practical recommendations.

<sup>2</sup> There are also criticisms of effective altruists' other methods, such as the Quality-Adjusted or Disability-Adjusted Life Years (Herzog 2016). Due to space constraints, I will not discuss them here. See (Karnofsky 2010; Todd 2015; Weathers 2016).

the effective altruism movement. Halstead et al enumerate several effective altruist organisations – spread across a wide range of causes – to establish this claim. For instance, the Future of Humanity Institute and the Machine Intelligence Research Institute aims to reduce the risks of human extinction; Animal Charity Evaluators and Sentience Politics aim to reduce animal suffering; 80,000 Hours and the Center for Applied Rationality coach individuals to do the most good, especially through their choice of careers (2017: 5-6). Each of these effective altruist organisations may endorse specific commitments – about which moral view to adopt, which values to promote, and which method to use – rendering their effective altruism thicker. Yet thin effective altruism remains in principle ecumenical. This point about ecumenicity is also supported by Peter Singer (2015: 79), Jeff McMahan (2017: 92), and William MacAskill (2017a, 2017b).

I take these recent statements seriously. My discussions will thus focus on thin effective altruism. One payoff of this focus is that my discussions refer to commitments that all effective altruists hold, and is thus relevant to all of them.

Thin effective altruism has the following features:

- (i) *Most Effective*. Effective altruism recommends individuals to do the most good by using careful reasoning and good evidence;
- (ii) *Moral View Ecumenicity*. Effective altruism is ecumenical among different moral theories and views;
- (iii) *Cause Ecumenicity*. Effective altruism is ecumenical among different causes; and
- (iv) *Method Ecumenicity*. Effective altruism is ecumenical among different analytical tools and methods of collecting evidence.<sup>3</sup>

*Moral View Ecumenicity* allows effective altruists to affirm different sources and standards of value(s) beyond increasing or maximise individuals' well-being. Effective altruism is not necessarily wed to consequentialism or welfarism. In principle, the grounds of effective altruism are left open. Effective altruism can recognise that values such as justice, freedom, equality or knowledge – or those identified by other moral theories and views – are intrinsically good, and may also be promoted. Individuals holding these views are not required to give them up. More of what is regarded as valuable within morality can be taken into account.

Second, this supports *Cause Ecumenicity*. Effective altruists may freely and legitimately engage in activities that promote different values – depending on the moral views they begin with.<sup>4</sup> Thus, MacAskill writes that 'decisions about cause selection involve *value judgement*' (2015a: 180; my emphasis). Similarly, the Centre for Effective Altruism and 80,000 Hours

---

<sup>3</sup> There is a worry about the relationship between *Most Effective* on the one hand, and the ecumenical moves on the other – the idea of 'doing the most good', which is a maximising notion, may not be comfortably incorporated into some non-consequentialist theories. If so, there may be a limit to how ecumenical effective altruism may be. I am, however, unable to pursue this meta-ethical issue further here.

<sup>4</sup> Some effective altruists, such as Halstead et al, claim that when it comes to global poverty and health, effective altruists' recommendations 'in fact would be endorsed by a range of moral theories' (Halstead et al 2017: 7). There are two ways of reading this. Endorsement could mean that the values promoted by these causes receive recognition within different moral theories, or that they receive prioritisation within those theories relative to other values. The former, while plausible, does not secure the further claim that individuals committed to different theories *should* contribute to the recommended causes. The promotion of other values may be regarded as having overriding weight. The latter secures the further claim, but at the expense of failing to take seriously the extent and depth of pluralism in individuals' views about the good and the right (Berlin 2013: 12-14).

ask individuals to consider their values and commitments when deciding which cause to contribute to.<sup>5</sup> *Moral View Ecumenicity* and *Cause Ecumenicity* thus allow effective altruists to engage in broader projects, beyond those which directly save or improve lives on narrowly welfarist terms – especially such as policy advocacy for institutional reform. At this point, effective altruists can already reject the criticism that effective altruism is, in principle, unable to account for, or justify engagement with, such causes.<sup>6</sup>

Of course, all this ecumenicity would amount to nothing if the tools used to measure the effectiveness of charities were not similarly ecumenical. That is, effective altruists' fixation with a method of limited applicability may – as earlier discussed – block their move to endorse or engage in a whole range of causes. *Method Ecumenicity* enters here. That effective altruists are now ecumenical to different methods of collecting evidence beyond RCTs, means that they may account for more effects in their evaluation of charities, other than the direct effects of small-scale interventions over a short time-span (Halstead et al 2017). For instance, the Open Philanthropy Project allows different methods to be used for collecting and assessing the evidence for long-term projects such as 'work aiming to influence policy or scientific research' (Karnofsky 2016). Thus, it is able to endorse engaging in such projects despite the fact that RCTs are inapplicable to them. And according to these different methods – which are better able to track indirect or long-term effects – policy advocacy may turn out to be very (if not the most) effective. In this way, *Method Ecumenicity* thus plays an enabling role in ensuring that *Moral View Ecumenicity* and *Cause Ecumenicity* are properly satisfied.

The ecumenical movies apply not only to individuals' donations to charities, but also to their career choices. On thick effective altruism, working in high-paying jobs – even those which are morally controversial – is recommended, because they are most effective. They generate the highest possible amount of income, which may then be donated to the most effective charities (MacAskill 2014). According to thin effective altruism, however, individuals may choose careers which are not most effective on thick effective altruist terms (and which would even be rejected on those terms). For instance, MacAskill argues that individuals may choose their careers based on the consideration of 'personal fit' (MacAskill 2015a: 148-155). He invites individuals to think about the following questions:

'How do I personally fit with this job? How satisfied will I be in this job? Am I excited by the job? Do I think I could stick with it for a significant period of time? How good am I, or could I become, at this type of work, compared to other people and compared to other careers I might choose?' (2015a: 148)

In answering these questions, an individual may judge that she is most suited for a career which is not the highest-paying. Thin effective altruism may readily accommodate this judgement. Because of *Moral View Ecumenicity*, individuals may begin from a plurality of moral views and values, reach different evaluations about which values to promote, and about how well different careers fit them. Effective altruists no longer need to take the hard-line, by insisting that an individual *has to* take the highest-paying job, even if it conflicts with her

---

<sup>5</sup> See the Cause Prioritization Tool (<https://www.effectivealtruism.org/cause-prioritization-tool/>) and Decision-making tool (<https://80000hours.org/career-decision/>), respectively.

<sup>6</sup> Here, I set aside Halstead et al's argument that being ecumenical towards different causes requires that individuals (including effective altruists) support charities 'insofar as they produce the most good impartially conceived, and do not support them for agent-relative reasons' (2017: 26). This argument violates *Moral View Ecumenicity*, insofar as the directive to support causes that produce the most impartial good for agent-neutral reasons is not one which may be readily endorsed by different moral views.

values.<sup>7</sup> Effective altruists may additionally appeal to *Cause Ecumenicity* to affirm that the careers chosen, despite not generating the highest amount of money, may nevertheless contribute to the promotion and maximisation of other values. Finally, they may also cite *Method Ecumenicity*, in allowing for individuals' career choices to be effective – when considered using alternative methods of collecting and assessing evidence, or when considered over a much longer time-span. This would also support a claim that certain “enabling” careers – such as Singer's or MacAskill's engagement in research and teaching – may have indirect effects which may, when taken cumulatively and over time, *ex post* lead us to regard the career as all-things-considered effective (Singer 2015: 55-56). And the fact that current methods do not (yet) allow us to easily determine such effects of these careers, need not preclude individuals from taking them up, *ex ante*.

### 3 Misguidedness and disagreement

The ecumenical moves allow individuals with different moral views, and who are committed to causes promoting different values, to be considered effective altruists. What is crucial is simply that they satisfy *Most Effective*. In this section I argue altruists' criticisms of purportedly less effective charities are inconsistent with their commitment to ecumenicity. Instead of being misguided, these charities (and the individuals who support them) may simply be in disagreement with effective altruists about which moral view to begin with, which values to promote, and which method is most appropriate. This is a conclusion that effective altruists should take seriously, in making their subsequent criticisms.

Consider effective altruists' criticism of individuals donating to purportedly less-effective charities – for instance, those that train guide dogs for the blind (Singer 2015: 110; MacAskill 2015a: 39), or to the Make-A-Wish Foundation (Singer 2015: 5-6), which satisfies the (often-fantastical) wishes of children with life-threatening medical conditions. The same criticism is levelled against these two fairly different charities – they do not make the most effective use of money. For instance, it takes on average around US\$40,000 to US\$50,000 to train and provide a guide dog for a single blind person (Singer 2015: 110; MacAskill 2015a: 39), and around US\$7,500 to make a single child's wish come true (Singer 2015: 6). In contrast, it takes only around US\$20 to US\$100 to prevent someone in a developing country from developing trachoma, which causes blindness (Singer 2015: 111), and around US\$3,400 to save a life (MacAskill 2015a: 54). Much less is required to significantly increase life expectancy through other measures such as providing de-worming services, anti-retroviral therapy, or insecticide-treated bed nets (MacAskill 2015a: 51-53). The challenge issued to individuals who wish to donate to charities that train guide dogs, or fulfil sick children's wishes, then, is whether doing so is ‘the *best* use of money’ (MacAskill 2015a: 51; his emphasis). Call this criticism *Less Effective*.

Effective altruists offer some explanations for why individuals donate ineffectively. Drawing on work by psychologists, Singer observes that many individuals who give small amounts of money to charities are ‘not so interested in whether what they are doing helps others’ (Singer 2015: 5). Instead, their knowledge that they are doing something altruistic ‘makes them feel

---

<sup>7</sup> This has exactly happened. 80,000 Hours now thinks that ‘only a small proportion of people’ *should* take up high-paying jobs in order to donate (MacAskill 2015b). They also recommend that individuals ‘don't take a career for the greater good if that career directly causes significant harm’ (Todd and MacAskill 2017). Among the five reasons provided, one of them centres on how the career relates to individuals' life plans and enjoyment – that is, on personal fit.

good, regardless of the impact of their donation’ (2015: 5). Psychologists describe these as ‘warm glow givers’. On Singer’s descriptions, warm glow givers are not particularly reflective. Their donations are based on ‘whatever cause tugs most strongly at their heartstrings’ (2015: 6), rather than on the evidence available concerning the effectiveness of the charities to which they donate. Their donations are guided by the wrong reasons, and made without the full light of evidence.

Another explanation is that individuals often donate to causes that they have some personal connection or attachment to. For instance, someone who has a family member die from cancer finds it ‘natural to want to direct [her] energies to fighting cancer’ (MacAskill 2015a: 40-41). Someone who is passionate about photographing nature may donate to charities that protect national parks. Someone who loves dogs may donate to dog shelters. Or, even more generally, people donate to charities that help only disadvantaged members of their society, rather than those living elsewhere (Singer 2015: 86). According to Singer, these are ‘commonly expressed dispositions and affections that effective altruists would consider *misguided grounds* for giving’ (Singer 2015: 86; my emphases). Again, people donate unreflectively, and thus misguidedly.

Consider *Less Effective* in light of *Moral View Ecumenicity* and *Cause Ecumenicity*, which make room for individuals to support causes that promote other values than those which are identified by consequentialist or welfarist moral views. The first observation is that different *types* of causes are being compared. Second, the values that are promoted by the causes being compared appear to be different. It is plausible to think that the values that are promoted by training guide dogs for blind people in our society, differs from those promoted by preventing others in another society from developing trachoma. For instance, the former may be aimed at promoting a fair and inclusive society. Such a society would be one in which the disabilities that an individual is born with (or ends up possessing) – in this case, blindness – often due to factors beyond her control, do not significantly and adversely affect her access to physical and social spaces, and, more generally, to the opportunities needed to pursue her life plans. Societies better promote the values of fairness and inclusiveness when they provide more accommodations for the disabled (Wolff 2009). Providing guide dogs for the blind is part of such accommodation. This is not the same as the aim of preventing individuals in other societies from developing trachoma and becoming blind. There, the aim appears to be to reduce the loss of well-being that accompanies loss of sight, without taking the location of, or our relationships with, the potential sufferers as morally relevant. The values of fairness and inclusiveness are not (completely and in all cases) reducible to that of reducing loss of well-being. For instance, a society that best promotes the former may not best promote the latter. As we have seen, much more resources are required to enact accommodations for people with disabilities, than to prevent others in other societies from becoming disabled in the same way. A similar analysis may be provided in comparing making the wishes of ill children in our society come true, and saving the lives of other children in other societies.<sup>8</sup>

If this is so, then Singer’s claims violate *Moral View Ecumenicity* and *Cause Ecumenicity*. Recall that these two ecumenical moves allow individuals to hold different moral views and support different causes. Yet as we have seen, the criticism of certain causes, such as providing guide dogs and fulfilling children’s wishes, is reached via evaluating these causes in terms of how well they promote the values prioritised by another moral view (which, in

---

<sup>8</sup> Alida Liberman makes a similar observation, in the context of the values promoted by Christians who donate to church activities (such as evangelism or worship), and those promoted by causes recommended by effective altruists (2017).

this case, appears to be welfarism). Of course, it is not that welfarism cannot be appropriately employed to evaluate those causes. It is rather that the moral views which would *best support* engagement with those causes have not been considered at all. This neglect conveys the idea that those views are not legitimate or acceptable as views on the basis of which we may evaluate the effectiveness of causes. This assumes that welfarism is the only acceptable view (at least in this context), and this in turn violates *Moral View Ecumenicity*. And insofar as this rules out these causes from being regarded as effective – when they in fact could be effective on different terms – *Cause Ecumenicity* is also violated.

Effective altruists, then, face the burden of explaining how *Less Effective* is consistent with *Moral View Ecumenicity* and *Cause Ecumenicity*. Such explanation is not foreclosed to them. They may argue that the different values which are promoted by the charities they criticise, *in fact* reduce to those promoted by the charities they recommend. Perhaps this may be successfully done, by drawing upon the discussions on “consequentialising” – representing the verdicts of non-consequentialist theories in consequentialist terms (see Portmore 2009; Brown 2011). That is, they could attempt to “welfarise” the values which are promoted by the charities that they criticise, in order to show that the charities they recommend best promotes those values instead. It is also open to effective altruists to argue that even though these individuals may begin from different views and values, they are *nevertheless* led to effective altruists’ conclusions about which charities to donate to – the promotion of non-welfarist values may turn out to be best served by the charities which effective altruists recommend. It may turn out that donating to the charities which effective altruists recommend, rather than those which train guide dogs, best promote the values of fairness and inclusiveness. What is crucial, however, is that effective altruists cannot assume, while the task is underway, that the results of those discussions will come out in their favour.<sup>9</sup>

Peter Singer appears to provide an argument of the first type. In responding to Melissa Berman’s argument that there are no precise answers to comparing charities that engage in different causes – such as preventing blindness, feeding the starving, rescuing animals, preventing rape, keeping glaciers frozen, or providing education or housing (Berman 2013) – Singer appeals to Quality Adjusted Life Years (QALYs). One QALY is equivalent to one year in without burdens (typically from diseases or disabilities). Different causes may then be evaluated on the basis of how much money they spend on gaining one QALY for a single individual. On this basis, ‘it is not, in principle, impossible to compare the benefits’ achieved by different charities (Singer 2015: 136). However, Singer’s response here violates *Method Ecumenicity*. Charities that try to keep glaciers frozen, or provide education or housing are not appropriately evaluated on the basis of QALYs, which are more suited for assessing charities that engage in direct health interventions such as preventing blindness, or feeding the starving (Herzog 2016). And when it comes to the very claim that needs argument – that the values which are enumerated and prioritised by different moral theories in fact reduce to those of (some form of) welfarism – Singer simply asserts that ‘all these charities really have a common goal: trying to improve the well-being of the poor’ (2015: 136).

In the meantime, and without assuming that the results will show the effective altruists to be wrong, I suggest that effective altruists should acknowledge that *Less Effective* is inconsistent with their commitment to ecumenicity. That is, it is presented in the absence of qualifications about the moral view being employed, or the values being considered. Indeed,

---

<sup>9</sup> My suggestions do not presuppose that the values that individuals with different moral views seek to promote are *in principle* incomparable. I am, however, unable to address the meta-ethical issue here. See (Hsieh 2016) for further discussions.

it appears as though the criticism is made from a perspective that all accept, concerning values which all agree should be promoted. *Less Effective*, then, does not properly acknowledge the fact that individuals holding different moral views may have different ideas about which values should be promoted, and thus different ideas about which causes they should support.

I am not claiming that the ecumenical moves render effective altruists unable to criticise or recommend charities. Instead, effective altruists need to pay more attention to the roles played by different moral views, and commitments to promoting different values, to their evaluations of a charity's effectiveness. I suggest that effective altruists' comparisons and criticisms of charities should ideally be accompanied by several qualifications, each corresponding to the ecumenical moves earlier described. Criticisms should be accompanied by explicit statements of the moral view which is used as the basis of evaluation, and the method of evaluation.. They should also be accompanied by detailed descriptions of the causes which they think the charities are aiming to further. Such descriptions should be specific rather than general, so that the values that are promoted by the cause are rendered clear. Consider the earlier comparison between charities which train guide dogs for the blind in our society, and those which prevent trachoma and blindness in another less-developed countries. This comparison appeared to be appropriate only given a very general description – perhaps something like “being about blindness”. We have seen that this general description masks the many differences (in terms of the values being promoted) between the charities. Being specific in describing the charities and causes avoids giving the mistaken impression that these differences are not morally significant. In any case, even if they are not significant, the issue needs to be settled by argument rather than avoidance.

Given this, *Less Effective* should ideally be restated thus: one value that we ought to promote is well-being and charities working to prevent trachoma and blindness in less-developed countries are more effective at promoting well-being, than charities which train guide dogs for the blind. Once the criticism is restated in this way, however, individuals who donate to their preferred charities (which are not recommended by effective altruists) have an immediate reply. They may easily concur with the effective altruists' criticism, yet deny that it motivates or justifies changing the charities they support. They may simply point out that they begin from a different moral view from that endorsed by effective altruists. Such a view provides them with a different starting set of values which they regard as the most important to promote, and which constitutes *for them* the relevant standard against which charities are evaluated. They then may donate to charities which further causes which are different from those recommended by effective altruists. And because of *Moral View Ecumenicity* and *Cause Ecumenicity*, they may do so.

At this point, we see that the force of *Less Effective* has been blunted. While these individuals may not be donating on the basis of evidence about the effectiveness of different charities, they may nevertheless be responding to good (rather than misguided) reasons given to them by their moral views. Upon reflection, they may simply conclude that they are in disagreement with effective altruists' starting moral views, or about which values ought to be promoted.<sup>10</sup> The unreflectiveness which effective altruists attribute to these individuals when they donate to charities which tug most strongly at their heartstrings, should not be taken as indicating their misguidedness. What may seem to be unreflectiveness may instead be

---

<sup>10</sup> Indeed, there are also psychological studies that suggest that individuals' donation decisions remain guided by their subjective preferences, which include the values they are committed to, even when presented with considerations of effectiveness (Berman et al 2018).

individuals' responses to the promotion of other values, which are conditioned or influenced by their endorsement of different moral views.<sup>11</sup>

The suggestion that effective altruists' criticisms and recommendations should be accompanied by explicit statements about the methods employed to assess the charities, is much more obvious. Indeed, effective altruists (especially effective altruistic organisations) have tended to be very careful with stating their methods explicitly. Such statements become even more important upon the introduction of *Method Ecumenicity*. As earlier discussed, effective altruists already recognise that different methods are needed to evaluate the effectiveness of different causes. For instance, while RCTs are not appropriate for evaluating the effectiveness of charities that engage in direct and local health interventions, they are less appropriate for evaluate the effectiveness of charities that engage in policy advocacy, or mitigating existential risk. Once we move beyond the controlled environments in which RCTs are appropriate, we face (at least) two additional dimensions of the evaluation of charities – “knock-on” effects beyond the direct effects of a charitable intervention,<sup>12</sup> and the time-span in which they occur.

Measuring knock-on effects is exponentially more complex than measuring the direct effects of a local and controlled intervention. This is partly because of how difficult it is to track *all* the effects of *every* action as the relevant time-span increases. A direct action may lead to several further effects, each of which would lead to several others, and so on. At each step along the way, the number of causally-related effects grows exponentially. This complexity is compounded by the fact that a significant portion of these effects may not be intended or even foreseeable, and thus even more difficult to identify and measure. As has been pointed out, effective altruists ‘just like everyone else, cannot possibly include estimates of *all the consequences* of [any particular intervention], from now *until the end of time*’ (Greaves 2017: 324; my emphases). The longer the time-span which is taken as relevant for the evaluation of charities' effectiveness, the more difficult it is to keep track of, and account for, all the effects of their activities.

The introduction of these dimensions, then, creates two additional tasks for effective altruists. First, they have to explain how we should select causally-related knock-on effects of a particular intervention. Here we are supposing that we cannot account for *all* – but merely a portion – of them for the purposes of evaluating the effectiveness of charities' activities. The question then arises as to *which sub-set of these effects we should take into account*, and how we should reach such a decision. Here, *Method Ecumenicity* complicates matters. Not only do different methods have different scopes of application, they may also regard different variables as salient. For instance, a quantitative method may diverge from a qualitative one, in terms of which effects they count as constituting evidence.

---

<sup>11</sup> We may supplement this with a different argument by Violetta Ionescu, who argues that ‘effective altruism is at most one part of ethical life’. The emotional responses of people who donate to charities that pull at their heartstrings are also important parts of an ethical life, for they are necessary for us to ‘develop deep relationships and live meaningful lives’ (2016: 148). This means that on some moral views, donating to charities based on emotional responses may be directly justified (even on effective altruist terms), without the need to provide additional arguments for how they may be indicative of responses to the promotion of other values. The discussions of effective altruism's relationship to ethical life more generally is tangential to my current concerns, and I set it aside.

<sup>12</sup> I adopt Hilary Greaves' (2017) terminology for such effects. Others within the effective altruism community sometimes refer to these as “flow-through” effects (Hurford 2016).

Consider this in the context of measuring the effects of an intervention, on individuals' lives. One method of gathering evidence is based on evaluating Well-being Adjusted Life Years, or WALYs (MacAskill 2015a: 39-40; Todd 2015). WALYs are collected by surveying, then aggregating, individuals' report of their subjective well-being across different social and health conditions. With these aggregated reports, we can make claims such as that typically 'doubling someone's income gives a 5-percentage increase in reported subjective well-being' (2015: 39). Such a method would track the subjective effects of well-being of any particular intervention. However, there are other methods of measuring the effects of an intervention. For instance, consider methods that measure how many objective constituents of a flourishing life are present.<sup>13</sup> Such methods may track other considerations such as whether the individuals concerned are subject to certain dominating relations with others, whether they lack affiliations with others within the society, whether they occupy positions which render them susceptible to certain kinds of risks (Wolff & de-Shalit 2007: 108-118), or whether they possess certain crucial functionings or basic capabilities (Sen 1992: 39-40). These other methods may allow us to say of certain individuals – notably, women in patriarchal societies, lower-caste members in hierarchical ones, or disabled individuals who have adapted to unaccommodating societies – that they in fact have more impoverished lives than what may be indicated by their subjective self-reports. This is because these methods may track other effects of the same intervention, beyond those which affect individuals' subjectively reported well-being. The specifics of these methods are not important for now. What is crucial is instead the general point that different methods will track and measure different effects of any particular intervention.<sup>14</sup>

Second, effective altruists have to explain *when to stop accounting for the knock-on effects* of any particular intervention, in evaluating the effectiveness of charities. Consider this in the context of Singer's discussion of MacAskill's career choice. In defending MacAskill's career choice, Singer appeals to the potential knock-on effects of being a teacher – being able to clarify students' thoughts on altruism, which would presumably motivate them to join the effective altruism movement. This is not idle speculation about the possible effects of MacAskill's career choice. While Singer probably does not intend to be self-congratulatory, we may find evidence for his claim about such knock-on effects, in the anecdote with which he begins his book – about one of *his* students, who took up a finance career and now contributes significantly to effective charities (Singer 2015: 3-4, 10). This lends support to, and renders more plausible, his claim that MacAskill's being a teacher may lead to knock-on effects that allow us to say of him that he has 'done more good than if he had gone into finance himself' (Singer 2015: 55).<sup>15</sup> *Method Ecumenicity* allows for this. Different actions (donating, earning to give, teaching, or policy advocacy) may all meet the demands of *Most Effective*, even though their intended effects are further away in the future.

Effective altruists are aware of, and are actively seeking solutions to, this issue of accounting for knock-on effects (Hurford 2016). But since there is no way of accounting for *all* the effects of any particular intervention, the complexity cannot be entirely overcome. Of course, effective altruists may nevertheless do well, by devising a unifying method which accounts for all (or most) of the knock-on effects which all other *available* methods are able to.

---

<sup>13</sup> The choice of method may also, at some level, be influenced by the moral view one begins with.

<sup>14</sup> The usage of different methods may even generate wildly different evaluations of the quality and effectiveness of *the same charity*. For instance, in 2013 the Wounded Warrior Project (a charity for wounded US veterans) received low, middle, or high ratings for its work depending on the evaluators (Cohen 2013).

<sup>15</sup> I am, of course, simplifying matters. MacAskill donates a significant portion of his income to effective charities – a consideration which has to be accounted for in the evaluation of his effectiveness.

Whether they in fact succeed, however, will depend on how amenable the plurality of methods are to unification – a discussion beyond the scope of this paper.

Let us now return to *Less Effective*, which was based on comparing charities which train guide dogs for the blind, and others which prevent trachoma. It should be clearer by now that the method for assessing the effectiveness of the latter (which turns on measuring how much it costs to prevent one instance of blindness) is inappropriate when applied to the evaluation of the effectiveness of the former. As with charities that engage in policy advocacy, the actions of charities which train guide dogs may take a much long time to pay off. The evaluations of their effectiveness is also more complicated. Indeed, it is difficult to say how many guide dogs need to be provided before we can say that blind people in our society are appropriately accommodated, and even more difficult to say when the values of fairness and inclusion are realised. The same may be said for the relationship between making a single policy change, and the point at which we say that justice (for immigrants, minorities, etc) has been achieved. *Less Effective* is not attentive to the demand of *Method Ecumenicity* – it compares the effectiveness of the direct interventions of both charities, without taking into account the appropriateness of the methods used to reach those evaluations.

In sum, individuals who donate to charities that effective altruists criticise (or do not recommend) may not be acting on the basis of wrong reasons – they may not be misguided. Criticisms of the charities they donate to have to be sensitive to the different moral views that these individuals adopt, the values they are trying to promote, and the methods that are used to evaluate the effectiveness of the charities' interventions.

Has my rehabilitation of individuals' donation patterns gone too far? It may seem that on my account, there is no room to criticise individual donors. This impression is mistaken. First, the fact that individuals may, on the basis of their commitment to different moral views and values, donate to these different charities, does not mean that they *should*. The effective altruists' may still appeal to *Most Effective*, according to which the charities which individuals donate to must be shown to be the most effective. Consider a person who donates to a charity which trains guide dogs for blind people in her society because she wishes to promote the value of fairness and equality (as described above). When she is presented with information that another charity that does exactly the same work in the same society is more effective, it seems that she should rightly redirect her donations to the more effective charity (Pummer 2016). So, whether these individuals' support of different charities meets the demands of *Most Effective*, depends on whether convincing evidence can be provided to support their claim that their donations are very (if not the most) effective in promoting *their chosen values* through their causes. Of course, this is a much more constrained criticism than *Less Effective*. But it has the benefit of being consistent with their commitment to ecumenicity. Second, effective altruists may also – as suggested earlier – argue that even though individuals begin from different views and values, they are nevertheless led back to effective altruists' conclusions about which charities to donate to. The question of whether these projects are fruitful, however, is one which I cannot address here.

#### **4 Thick and thin**

Recall that Halstead et al draw the distinction between the idea of effective altruism, and the actions and recommendations of effective altruists. This is the distinction between thin and thick effective altruism. Effective altruism is, in principle, thin. Individuals who consider

themselves effective altruists may, however, be committed to specific moral views and to the promotion of specific values – they are thick effective altruists. We get from thin to thick effective altruism by “plugging in” additional commitments to specific moral views and values. With this distinction, proponents of effective altruism may advance the following objections, both of which begin with the observation that my discussions in Section 3 are about the criticisms of effective altruists – in this case, Peter Singer and William MacAskill – rather than effective altruism. First, the fact that these thick effective altruists are not properly ecumenical, does not mean that thin effective altruism is not so. Arguments against the actions and recommendations of effective altruists are ‘not necessarily valid against the idea of effective altruism as expressed by its definition’ (Halstead et al 2017: 3-4). Thin effective altruism remains ecumenical, even if thick effective altruists are not. The second objection is resolute: thick effective altruists need not be ecumenical. Thin effective altruists’ commitment to ecumenicity does not extend to thick effective altruism. Criticising thick effective altruism for not being ecumenical, would thus be to mistake the relationship between thin and thick effective altruism.

The first objection invites an observation and a question. If it is true that Singer and MacAskill are indeed thick effective altruists, we will have to reconsider the status of the criticisms that they present. Their criticisms (of charities which train guide dogs for the blind, or fulfil children’s wishes) are presented in books introducing the *idea* of effective altruism. That this is so, gives the impression that they are intended as part of the elaborations of the ‘very simple idea’ of effective altruism itself. On the basis of the distinction that Halstead et al draw, such an impression would be mistaken. The criticisms and conclusions they make *cannot* be drawn from the idea of effective altruism. Instead, they are criticisms grounded in the specific moral views and commitments that Singer and MacAskill hold.<sup>16</sup> These views and commitments would then have to be made explicit, rather than hidden. Their books are not just about the idea of effective altruism, but more accurately about a specific variant of thick effective altruism.

The question is about whether any moral view or cause may be plugged in to thin effective altruism. Consider moral views such as animism, asceticism, egoism, solipsism, or nihilism. Or consider charities which seek to increase donations to wealthy universities, to increase the reach of the arts within wealthy societies, to develop technologies that improve everyday electronics, or to improve the sophistication with which people appreciate alcoholic drinks. Could these views and causes be plugged in to thin effective altruism, giving rise to thick effective altruisms which attempt to be most effective at promoting the associated values? With what has been said so far, it appears that they could – all they need is to show that they meet the demands of *Most Effective*. The question, then, is how ecumenical thin effective altruism can (and want to) be. Here, proponents of effective altruism may take a lesson from how the term ecumenicity is sometimes used in Christianity – to refer to the idea that different denominations are nevertheless united in a worldwide communion. This ecumenicity is not boundless – it excludes groups which are clearly not Christian. Effective altruists may claim, likewise, that thin effective altruism is not boundlessly ecumenical. It excludes these peculiar moral views (some of which, such as egoism, may not even count as moral), and these causes. But this is a position that effective altruists will need to elaborate on. Specifically, they will need to clarify which *additional* commitments they endorse, which allow them to rule out these views and causes. They will also need to explain how these

---

<sup>16</sup> This helps us to make better sense of Gabriel’s cryptic statement that thick effective altruism ‘explain many of its judgments and capture much of what makes it unique’ (2017: 459).

additional commitments nevertheless leave enough room for a wide range of moral views and causes to be regarded as effectively altruistic. Effective altruists cannot simply stop at giving brief declarations about the ecumenicity of thin effective altruism. Much more will have to be said.

In response to the second objection, we can acknowledge that thick effective altruism is needed to do the practical work of critiquing and recommending charities. However, this does not fully address the worries about ecumenicity raised in the earlier section. Suppose we take this resolute position to indicate that the criticisms and recommendations of thick effective altruists are convincing only to those who are committed to the same moral views, and to the promotion of the same values (which are regarded as) identified by those views. If this is so, thick effective altruists would have to simply concede that there is nothing more to say to individuals who disagree with them about those commitments. Of course, such individuals may still be convinced by the conclusions of other thick effective altruists who begin with commitments that they share. However, it is also possible that such individuals may simply reject all forms of thick effective altruism. Given that we live in a world containing a wide range of moral views, and given the severity and urgency of the issues which effective altruists are trying to tackle, this may be too much of a cost to bear. There are potentially great benefits of ensuring that thick effective altruists' criticisms and recommendations are convincing to those holding on to a wide range of moral views. Even thick effective altruists, then, have good reason to try to be as ecumenical as they can.<sup>17</sup>

That even thick effective altruists should be ecumenical, is a position that effective altruists actually hold. In many of their discussions, there are often promises or assertions that their criticisms and recommendations would be endorsed by a range of moral views. The resolute stance that is described by the second objection, is one which does not do justice to effective altruists' concern with convincing and recruiting others who hold different views, and thus is a stance to which they are not committed. If so, the second objection does not dismiss the worries about ecumenicity raised in the earlier section.

## 5 Conclusion

Effective altruists are correct to acknowledge, and take seriously, the plurality in moral views and commitments among those to whom they address their criticisms and recommendations. Their statement of their ecumenicity is a welcome move. However, and as I have argued, their criticisms and recommendations of some charities and causes are not consistent with their commitment to ecumenicity. I have suggested, in light of this, that effective altruists should significantly revise their criticisms and recommendations. In order to fulfil the demands of ecumenicity, these criticisms and recommendations should ideally be accompanied by explicit statements about the moral view which is endorsed, the values which are taken as requiring promotion, and the methods used to evaluate effectiveness. This

---

<sup>17</sup> Peter Singer acknowledges that (thick) effective altruism is very demanding, but notes that when addressing a broad public audience, to advocate something very demanding is 'to risk putting most people off doing anything, and therefore likely to do less good than advocating a *lower* standard' (2016: 164; my emphasis). That he describes such accommodation as a *lowering* of standards, relates to his description elsewhere that people often acts on the basis of misguided reasons (2015: 86). Effective altruists may agree to accommodation of this kind, without agreeing with Singer's explanations. Instead, accommodation can be motivated by the demands of ecumenicity. That is, recommendations that leave no room for different values and moral views are unlikely to be convincing to, or endorsed by people with such commitments.

will mitigate their dismissiveness of individuals and charities that do not engage in activities that further the values that some thick effective altruists are committed to – and which has led many people to recoil from effective altruism in general.

### Acknowledgment

I am grateful for comments by Brian Berkey, Sebastian Bishop, Susanne Burri, Tom Dougherty, Jessica J.T. Fischer, Kailing Fu, Justin Holder, Sonny Kim, Jeff McMahan, and the anonymous reviewers for the journal. An earlier draft was presented at Applied Ethics Discussion Group at University of Oxford, and the Society for Applied Philosophy Annual Conference 2018 in Utrecht. I thank the participants for their comments.

### References

- Berman, Melissa. 2013. In charitable giving, no ‘hierarchy of goodness’. *The New York Times* (19 August). Available at: <https://www.nytimes.com/2013/08/20/opinion/in-charitable-giving-no-hierarchy-of-goodness.html>
- Berman, Jonathan et al. 2018. Impediments to effective altruism: The role of subjective preferences in charitable giving. *Psychological Science* 29(5): 834-844.
- Berlin, Isaiah. 2013. *The Crooked Timber of Humanity*, ed. Henry Hardy. Princeton: Princeton University Press.
- Brown, Campbell. 2011. Consequentialize this. *Ethics* 121: 749-771.
- Clough, Emily. 2015. Effective altruism’s blind spot. *Boston Review* (14 July). Available at: <http://bostonreview.net/world/emily-clough-effective-altruism-ngos>.
- Cohen, Rick. 2013. One charity, many different ratings: What’s a donor to do? *Nonprofit Quarterly* (8 January). Available at: <https://nonprofitquarterly.org/2013/01/08/one-charity-many-different-ratings-whats-a-donor-to-do/>
- Gabriel, Iason. 2017. Effective altruism and its critics. *Journal of Applied Philosophy* 34(4): 457-473.
- Gray, John. 2015. How & how not to be good. *New York Review of Books* (15 May). Available at: <http://www.nybooks.com/articles/2015/05/21/how-and-how-not-to-be-good/>.
- Greaves, Hilary. 2017. Cluelessness. *Proceedings of the Aristotelian Society* 116(3): 311-338.
- Halstead, John, et al. 2017. Effective altruism: An elucidation and a defence. *Centre for Effective Altruism Blog*. Available at: <https://www.centreforeffectivealtruism.org/blog/effective-altruism-an-elucidation-and-a-defence/>.
- Herzog, Lisa. 2016. Can ‘effective altruism’ really change the world? *OpenDemocracy* (22 February). Available at: <https://www.opendemocracy.net/transformation/lisa-herzog/can-effective-altruism-really-change-world>.
- Hsieh, Nieh-hê. 2016. Incommensurable values. The Stanford Encyclopedia of Philosophy (Spring 2016 Edition), ed. Edward N. Zalta. Available at: <https://plato.stanford.edu/archives/spr2016/entries/value-incommensurable>.
- Hurford, Peter. 2016. Five ways to handle flow-through effects. *Effective Altruism Forum* (28 July). Available at: [http://effective-altruism.com/ea/zt/five\\_ways\\_to\\_handle\\_flowthrough\\_effects/](http://effective-altruism.com/ea/zt/five_ways_to_handle_flowthrough_effects/).
- Igneski, Violetta. 2016. Living a meaningful and ethical life in the face of great need: Responding to Singer’s The Most Good You Can Do. *Journal of Global Ethics* 12(2): 147-153.
- Karnofsky, Holden. 2010. DALYs and disagreement. *Give Well Blog* (21 April). Available at: <https://blog.givewell.org/2008/08/22/dalys-and-disagreement/>.
- Karnofsky, Holden. 2016. Hits-based giving. *Open Philanthropy Project Blog* (4 April). Available at: <https://www.openphilanthropy.org/blog/hits-based-giving>.
- Krishna, Nakul. 2016. Add your own egg. *The Point Magazine* (Issue 11). Available at: <https://thepointmag.com/2016/examined-life/add-your-own-egg>.
- Lieberman, Alida. 2017. Effective altruism and charity: possibilities for productive collaboration. *Essays in Philosophy* 18(1): Article 2.

- MacAskill, William. 2014. Replaceability, career choice, and making a difference. *Ethical Theory and Moral Practice* 17: 269-283.
- MacAskill, William. 2015a. *Doing Good Better: Effective Altruism and How You Can Make A Difference*. New York: Gotham Books.
- MacAskill, William. 2015b. 80,000 Hours thinks that only a small proportion of people should earn to give long term. *80,000 Hours Blog*. Available at: <https://80000hours.org/2015/07/80000-hours-thinks-that-only-a-small-proportion-of-people-should-earn-to-give-long-term/>.
- MacAskill, William. 2017a. Effective altruism: Introduction. *Essays in Philosophy* 18(1): Article 1.
- MacAskill, William. 2017b. Introducing CEA's guiding principles. *Effective Altruism Forum* (8 March). Available at: [http://effective-altruism.com/ea/181/introducing\\_ceil\\_guiding\\_principles/](http://effective-altruism.com/ea/181/introducing_ceil_guiding_principles/).
- McMahan, Jeff. 2017. Philosophical critiques of effective altruism. *The Philosophers' Magazine* 73: 92-99.
- Nussbaum, Martha. 1997. If Oxfam ran the world. *London Review of Books* 19(17): 18-19.
- Portmore, Douglas. 2009. Consequentializing. *Philosophy Compass* 4(2): 329-347.
- Pummer, Theron. 2016. Whether and where to give. *Philosophy & Public Affairs* 44(1): 77-95.
- Rubenstein, Jennifer. 2016. The lessons of effective altruism. *Ethics & International Affairs* 30(4): 511-526.
- Sen, Amartya. 1992. *Inequality, Re-examined*. Oxford: Clarendon Press.
- Singer, Peter. 2015. *The Most Good You Can Do: How Effective Altruism Is Changing Ideas About Living Ethically*. New Haven: Yale University Press.
- Singer, Peter. 2016. The most good you can do: A response to the commentaries. *Journal of Global Ethics* 12(2): 161-169.
- Srinivasan, Amia. 2015. Stop the robot apocalypse. *London Review of Books* 37(18): 3-6.
- Todd, Benjamin. 2015. We care about WALYs not QALYs. *Effective Altruism Forum* (13 November). Available at: [http://effective-altruism.com/ea/pu/we\\_care\\_about\\_walys\\_not\\_qalys/](http://effective-altruism.com/ea/pu/we_care_about_walys_not_qalys/).
- Todd, Benjamin & William MacAskill. 2017. Is it ever okay to take a harmful job in order to do more good? An in-depth analysis. *80,000 Hours*. Available at: <https://80000hours.org/articles/harmful-career/>.
- Weathers, Scott. 2016. Can 'effective altruism' change the world? It already has. *OpenDemocracy* (29 February). Available at: <https://www.opendemocracy.net/transformation/scott-weathers/can-effective-altruism-change-world-it-already-has>.
- Williams, Bernard. 1973. A critique of utilitarianism. In *Utilitarianism: For and Against*, eds. J.J.C. Smart and Bernard Williams, 77-150. Cambridge: Cambridge University Press.
- Wolff, Jonathan and de-Shalit, Avner. 2007. *Disadvantage*. Oxford: Oxford University Press.
- Wolff, Jonathan. 2009. Disability among equals. In *Disability and Disadvantage*, ed. Kimberley Brownlee and Adam Cureton, 112-137. Oxford: Oxford University Press.