

Effective Altruism's Underspecification Problem

Travis Timmerman

1. Introduction

Whether any given act is supererogatory, obligatory, merely permissible, or impermissible depends upon the alternative acts available to the agent. But what exactly are the *relevant* alternative acts available to an agent? It turns out that this is a surprisingly difficult question to answer, yet it's one on which any complete normative ethical theory must take a stance. It's also one on which any effective altruist must take a stance. This may be unwelcome news for effective altruists since, as I will demonstrate, each of the dominant views in the literature generates verdicts that are (i) implausible in their own right and (ii) seemingly at odds with typical effective altruist commitments.

Considering a particular case will help make this issue less abstract:

The Gig Brandi has been invited by her friend, Chad, to attend his musical gig. Brandi can easily decide to attend the gig, and then decide at the gig to be supportive of Chad, which would be the best outcome. Unfortunately, Chad is a mediocre musician. Consequently, Brandi would not in fact decide to be supportive of Chad if she decided to attend his gig due to being irritated with Chad's performance—even though she *could* decide at the gig to be supportive. Since Chad would be deeply hurt, this would be the worst outcome. Brandi could alternatively decide not to attend Chad's gig, which would be better than the worst outcome, yet worse than the best outcome.¹

To be sure, Brandi *can* decide to attend the gig, and once, there she *can* decide to be supportive of Chad. However, suppose that just isn't what Brandi *would* do if she were to attend. Here's the tricky philosophical question. Is Brandi morally obligated to accept or decline the invitation? Roughly, possibilists hold that Brandi is obligated to accept because accepting is part of the best series of acts she can perform over the course of her life. By contrast, actualists hold that Brandi is obligated to decline because what would actually happen if Brandi declines is better than what would actually happen if she accepts.² So, unlike possibilists, actualists hold that facts about how an agent *would* freely act partly determine an agent's obligations.

What is recognizably today's actualist/possibilist debate can be traced back to Holly (Goldman) Smith's seminal paper, in which cases structurally identical to *The Gig* originated.³

¹ This example is drawn from Cohen and Timmerman (2016, p. 1).

² Given my case and the colloquial description of actualism and possibilism, this debate may appear to only present problems for maximizing consequentialist views. However, the issues raised by the actualism/possibilism debate apply to a wide range of normative ethical views. For instance, just with respect to consequentialism, this debate applies to satisficing consequentialist views, where the threshold for permissibility is determined by the goodness of the act relative to the agent's available act-alternatives. In order to determine the relative goodness of any given act, one must determine what the relevant act-alternatives are, and that requires taking a stance on the actualism/possibilism debate. More generally, this debate also applies to normative ethical views which hold that, all else being equal, one is obligated to maximize the good. Again, one would need to know which acts are the relevant options for the agent in order to determine which act-set the agent is obligated to perform in such cases. Since the actualism/possibilism debate concerns identifying the scope of acts that are *relevant options* for the agent, it seems to me that every normative ethical view must take a stance on the debate.

³ Goldman (1976).

Goldman's paper helped illustrate that any plausible normative ethical theory has to take a stance on the scope of the options available to the agent. Since the deontic status of an act depends upon the relevant options available to the agent, a normative ethical theory will have to take a stance on what the relevant options are and, of course, that requires taking a stance on the actualist/possibilist debate.

As William MacAskill's contribution to this volume illustrates, the term *effective altruism* has been understood in different, yet related, ways. It may be understood as a normative claim that one should do the most good they can.⁴ It may also be understood as a non-normative project, whereby one is committed to doing the most good they can with some or all of their expendable resources, irrespective of whether they believe they're obligated to act in this way.⁵ In this volume, MacAskill defines it thusly as "the use of evidence and reason to figure out how to benefit others as much as possible, and the taking of action on that basis."⁶

In this chapter, I'll understand *effective altruism* broadly to include both the normative thesis and non-normative project conceptions. My general aim is to argue that any understanding of effective altruism (either as a normative thesis or as a non-normative project) that does not take a stance on the actualism/possibilism debate is woefully incomplete. I will be interested in what pursuing the effective altruist project amounts to for those committed to actualism or to possibilism and how pursuing such a project maps onto the current practices and recommendations of those who currently self-identify as effective altruists. I will also argue that effective altruists should adopt my favored view, viz. hybridism, which avoids the independent problems to which actualism and possibilism are subject. I also argue that hybridism better coheres with each conception of effective altruism than its alternatives. Much of what I write applies to each existing conception of effective altruism. When that isn't the case, I will specify the particular conception I am invoking.

This chapter is structured as follows. In the next section, I will provide an overview of the actualism/possibilism debate and, in doing so, offer a brief defense of my positive view, viz. hybridism. After that, I illustrate how this debate bears directly on effective altruism literature. I then argue that effective altruism is subject to a dilemma. Effective altruists', at times, implicit actualist assumptions (*i*) commit them to conclusions seemingly antithetical to what typical effective altruists actually believe, as well as the spirit of the movement and (*ii*) undermine effective altruists' arguments against moral offsetting and giving to charities close to the heart. Yet, effective altruists', at times, implicit possibilist assumptions (*iii*) also commit them to conclusions seemingly antithetical to what typical effective altruists actually believe, as well as the spirit of the movement and (*iv*) undermine typical responses to demandingness worries for the normative conception of effective altruism.⁷ I argue that the best way out of the dilemma is to accept hybridism.

2. Actualism, possibilism, and hybridism

Actualism and possibilism may be defined more precisely as follows:

⁴ Singer (2016, p. vii).

⁵ Karnofsky (2013); MacAskill (2015).

⁶ Effective altruism does not officially take a stance on the good or even whether people are morally obligated to maximize the good. So, effective altruism does not assume impartial consequentialism, even though some prominent effective altruists are utilitarian. Effective altruists may, and some do, adopt the evidential reason component of effective altruism, while remaining skeptical about the maximizing component.

⁷ For a defense of such a worry, see Timmerman (2015a).

Actualism: At t an agent S morally ought to ϕ at t' iff (1) ϕ -ing at t' is an act-set under S 's control at t , and (2) what would happen if S were to ϕ at t' is better than what would happen if S were to perform any incompatible maximally specific act-set under S 's control at t .⁸

Possibilism: At t an agent S is obligated to ϕ at t' iff ϕ -ing at t' is part of the best act-set that S can perform from t to the last moment that S can possibly perform an act.

As *The Gig* suggests, possibilism is generally much more demanding of imperfect agents than actualism since possibilism requires agents to do the most good they can over the course of their entire life.⁹ Actualism, however, doesn't require nearly as much from imperfect agents. This is because actualism allows agents to avoid incurring obligations to do good in certain cases in which the agent would freely choose not to bring about the good. In *The Gig*, for instance, at the time Brandi received the invitation, it is not under her control to <attend the gig and be supportive of Chad>. Since she would freely choose *not* to be supportive of Chad if she chooses to attend the gig, Brandi avoids incurring an obligation to <attend the gig>. Possibilism, however, does not let Brandi off the hook simply because she is disposed to act wrongly. Possibilism, but not actualism, entails that unless an agent does the most good she can over the course of her entire life, she will have acted wrongly at some point in her life. In this respect, it's a more demanding view.

There are two standard objections to actualism in the literature. To help illustrate them, it will be beneficial to consider another case:

Selfish Sally: Sally has exactly \$500 in her account and had planned to use it to purchase non-refundable non-transferable tickets to see a Bad Religion concert tomorrow. Unless Sally spends the \$500 on herself today, she will face the following choice tomorrow: use the \$500 to give to an effective charity, which would result in three innocent lives being saved, or purchase the concert ticket. It's not presently under Sally's control to both keep the money in her account today and use the money to save the three lives tomorrow. That is, she cannot ensure today that she uses the money to save the three lives tomorrow. However, it is presently under her control to keep the money in her account today. Moreover, if she does keep the money in her account today it would be under her control tomorrow to use the money to save the three lives. At that point, all she has to do to ensure that she saves the three lives is to intend to use the \$500 to save the three lives. That is just not what she would do if she finds herself in that situation. Finally, Sally is

⁸ Unlike this formulation of actualism, early formulations did not build in a control condition. See Goldman (1976); Sobel (1976); Jackson and Pargetter (1986). These versions of actualism are subject to devastating problems. Most notably, they violate the principle of normative inheritance Portmore (forthcoming, ch. 4) and they generate conflicting obligations without saying which obligation takes priority Cohen and Timmerman (2016, pp. 11-12); Kiesewetter (2015, pp. 929-934); Portmore (2011, pp. 181-183). Subsequent versions of actualism built in a control condition, avoiding this problem. See, for instance, Goldman (1978, p. 202); Bykvist (2002, pp. 61-64); Jackson (2014). Douglas Portmore's (2011) and Jacob Ross's (2012) securitist views also count as versions of actualism for the purposes of this paper.

⁹ There are exceptions, however. Actualism might require certain agents to make demanding sacrifices now (e.g. doing something unpleasant to develop a more altruistic character) in order to ensure that they do more good in the future. Possibilism, by contrast, wouldn't require such sacrifices. I thank Michelle Hutchinson for raising this issue. See Timmerman and Swenson (forthcoming) for discussion of other ways actualism can be more demanding than possibilism.

aware of these facts and consequently decides to purchase the concert ticket for herself today rather than tomorrow.¹⁰

Actualism entails that there is no time at which Sally has an obligation to save the three lives tomorrow. This is because actualism does not even regard saving the three lives as a relevant *option* for Sally in spite of the fact that actualists grant that this is something Sally *can* do. She can save the three lives by simply choosing to save the money today and, once tomorrow rolls around, choosing to use the money to save the three lives. The problem is that actualism seems to get Sally off the hook too easily. In other words, it's not demanding enough:

The Not Demanding Enough Objection: Actualism permits an agent *S* to avoid incurring any moral obligation to ϕ , which *S* can easily fulfill, simply in virtue of *S*'s rotten moral character.¹¹

To illustrate the second objection to actualism, suppose that if Sally purchases the concert tickets today, the deaths of the three people would be moderately painful, whereas if she purchases the concert tickets tomorrow their deaths would be extremely painful. According to actualism, then, Sally ought to purchase the concert tickets today, causing three people to suffer a moderately painful death. This is because, of all the acts presently under her control, purchasing the tickets today would result in the best outcome. Moreover, Sally would presumably be immune from moral criticism since, according to actualism, she fulfilled her moral obligations and, we can suppose, did so for the right reasons. This consequence of actualism, however, seems absurd to many to since (actualists and possibilists agree) Sally *can* ensure that none of the three people are harmed at all. Simply stated, actualism prescribes bad behavior:

The Bad Behavior Objection: Actualism prescribes bad behavior, and acting on such prescriptions (according to actualism) presumably renders¹² an agent *S* immune from moral criticism, even when *S* can easily refrain from such behavior.¹³

2.2 Problems with possibilism

The aforementioned objections to actualism have led many to accept possibilism instead. However, possibilism is subject to at least one comparably difficult objection. As a result of ignoring facts about how agents would act in the future, possibilism generates action-guiding obligations that would result in the worst possible outcome. To illustrate, suppose that Sally is offered a deal where if she purchases the concert ticket today, then the profits from the ticket would be used to save two lives. Sally still can, of course, use the money to save three lives tomorrow if she has the money at that time. However, if she purchases the concert tickets tomorrow, then no lives would be saved. Possibilism entails that Sally ought to forgo the opportunity to purchase the tickets today in order to save two lives, even though this would result

¹⁰ This case is a modified version of a case given in Timmerman (2015b, pp. 1,516-1,517) and in Timmerman and Cohen (2016, p. 677).

¹¹ Jackson and Pargetter (1986, p. 240); Zimmerman (2006, p.156); Portmore (2011, p. 207); Baker (2012, pp. 642-3); Timmerman (2015b, pp. 1,512-1,513).

¹² This is not strictly entailed by actualism, but it is entailed by actualism coupled with widely accepted axioms about moral blameworthiness.

¹³ Wedgwood (2009); Ross (2013); Timmerman and Cohen (2016); Zimmerman (2017, p. 121).

in her purchasing the same tickets tomorrow and, consequently, failing to save anyone.¹⁴ Possibilism prescribes this even though, no matter what Sally intends to do today, she would freely choose to *not* save the three lives tomorrow. This objection may be formulated more precisely as follows:

The Worst Outcome Objection: Possibilism entails that an agent *S* can have an action-guiding obligation to ϕ even when ϕ -ing entails that *S* would perform an act-set that is deeply morally wrong and that is worse than the act-set *S* would perform if *S* were to $\sim\phi$.¹⁵

As these cases illustrate, while actualism and possibilism each enjoy some intuitive support, both generate intuitively incorrect verdicts in certain cases. I will now provide a *prima facie* defense of a hybrid view that, I believe, avoids the problems associated with actualism and possibilism.

2.3 Hybridism

Hybrid views posit two distinct moral “oughts”, one actualist in nature and one possibilist in nature. These oughts are meant to jointly track the insights of both actualism and possibilism, yet be immune from the three aforementioned objections. Given space limitations, I cannot provide a full defense of any particular hybrid view here, so my goal is to make a *prima facie* case for hybridism. My more important goal is to illustrate how issues that arise in the effective altruism literature hinge on the actualism/possibilism debate in ethics. Although I am making a *prima facie* case for hybridism, different versions of the view have already been developed and fully defended in the literature.¹⁶ In this section, I will focus on my favored version of hybridism, viz. *Single Obligation Hybridism* (SOH). In its simplest form, SOH posits a possibilist moral *obligation* that picks out the criterion of right and an actualist moral *ought* that functions as a sort of decision procedure. The actualist moral ought is an action-guiding practical ought, not a moral obligation. Hybridist views take agents’ shortcomings into its prescriptions, sometimes telling an agent to perform a wrong act now in order to prevent that person from performing an even worse act at a later time. SOH may be formulated more precisely as follows.¹⁷

¹⁴ If this is not already intuitive, imagine that the stakes are much higher. We can suppose that there are a billion and one lives at stake and that Sally has the opportunity to save one billion of those lives if she purchases the ticket today.

¹⁵ Goldman (1976, pp. 469-70); Woodard (2009, pp. 219-220); Portmore (2011, p. 211); Timmerman and Cohen (2016, p. 674).

¹⁶ See Timmerman (2015b) and Timmerman and Cohen (2016).

¹⁷ I am using the term *obligation* narrowly as shorthand for *moral obligation* and the term *ought* broadly to refer to any claim about how one should act within any normative domain. So, a moral obligation is one type of moral ought, while a practical moral ought is another type of moral ought. We could also speak of what one prudentially ought to do, what one legally ought to do, and so on. This formulation of SOH is a simplified version of the one given in Timmerman and Cohen (2016, pp. 682-683). The simpler version of SOH suffices for the purposes of this paper, however, since none of my arguments hinge on the issues addressed in the more complex definition.

Single Obligation Hybridism:

- (a) **Possibilist Moral Obligation:** At t an agent S has a possibilist moral obligation to ϕ at t' iff ϕ -ing at t' is part of the best series of acts that S can perform from t to the last moment that S can possibly perform an act.
- (b) **Actualist Practical Ought:** At t an agent S has most practical reason to ϕ at t' iff ϕ -ing at t' is under S 's control at t and ϕ -ing at t' is either (i) identical to the maximally-specific possibilist obligation that S has at t , (ii) a rationally permissible supererogatory act, or (iii) is the least rationally impermissible, all things considered, act-set presently under S 's control at t . There is an act-set that satisfies (iii) iff no act-set presently under S 's control at t satisfies conditions (i) or (ii).

The technical details of SOH are not centrally important for the purposes of this chapter. The most important elements of the view, for the purposes of this paper, may be understood by considering its applications in particular cases. Consider *Selfish Sally*. SOH entails that Sally has a possibilist moral obligation to refrain from purchasing the ticket today in order to use it to save the three lives tomorrow. This is because the possibilist obligation, which picks out the criterion of right, holds that agents are obligated to perform each act that is part of the best series of acts they can perform. The possibilist obligation can also serve to ground reactive attitudes. On plausible accounts of blameworthiness, Sally can be blameworthy to the extent she culpably fails to fulfill her possibilist obligations. Thus, SOH is immune from the *Not Demanding Enough Objection* and from the *Bad Behavior Objection*.

At the same time, SOH is immune from the *Worst Outcome Objection* because the actualist ought, not the possibilist obligation, is action-guiding. The actualist ought prescribes performing the act that would result in the best outcome *from among the set of acts presently under the agent's control*. Sometimes this is identical with the possibilist obligation and sometimes it isn't. This practical ought then serves the purpose of minimizing wrongdoing in light of one's present circumstances. So, in *Selfish Sally*, SOH entails that Sally practically ought to forgo saving the three lives tomorrow in favor of saving the two lives today. Hybridism tells Sally to perform a wrong act now (i.e. saving two instead of three innocent lives) in order to prevent herself from performing an even worse act later (i.e. saving no lives). Stated over simplistically, hybridism tells effective altruists to act like actualists, even though they are obligated to act like possibilists. Now that I have made my prima facie case for hybridism, I'll turn to the implications of this debate for effective altruism.

3. A dilemma: effective altruists' contradictory assumptions

The actualism/possibilism debate has, until now, been completely overlooked in the effective altruism literature. However, a number of people who self-identify as effective altruists implicitly appeal to actualist or possibilist considerations in their work. I'll use the term *effective altruist* to refer to someone who believes that they ought to be doing the most good they can either because they endorse effective altruism as a normative thesis or because they have adopted effective altruism as a non-normative project. So, as I am using the term, any effective altruist believes that they ought, in some sense, to be doing the most good they can. Those who accept the normative thesis believe they are *morally obligated* to do the most good they can.

Those who have only adopted effective altruism as a non-normative project believe they ought, qua effective altruist, to do the most good they can. This ought is not referring to a moral obligation. Such effective altruists have adopted the project of doing the most good they can because they believe such a project to be worthy of pursuit, likely in virtue of the fact that they perceive it as morally good (even supererogatory). In order to successfully achieve the aim of their project, they need to do the most good they can. So, the ought in question is one that picks out a necessary means to achieve a (morally important) aim. Analogously, someone who takes on the project of being vegan because they believe it to be a (non-obligatory) project worthy of pursuit ought, in this sense, not to consume animal products. Someone who takes on the project of writing a scientifically informed op-ed about vaccines ought, in this sense, to learn the relevant science. In short, this ought doesn't refer to a moral obligation; it refers to a necessary means of achieving one's goal.

Here's the problem. Effective altruists implicitly appeal to actualism when warding off concerns that effective altruism is too demanding, too impractical, or too out of sync with the "real world".¹⁸ Yet, they also implicitly appeal to possibilism when warding off concerns that effective altruism licenses bad behavior or lets some agents off the hook too easily. Since actualism and possibilism are defined in terms of moral obligations, and since they make incompatible claims about which possible outcomes are relevant options for the agent, they are contradictory. Thus, effective altruists cannot consistently appeal to both in their theorizing.

3.1 Effective altruists' implicit actualist assumptions

Two prominent effective altruists, Peter Singer and William MacAskill, frequently implicitly appeal to actualism in their work on effective altruism and specifically do so when trying to assuage concerns about the demandingness of being an effective altruist.¹⁹ The demandingness concerns mainly arise for effective altruism understood as a normative thesis (which is how Singer understands it), rather than understood as a non-normative project (which is how MacAskill understands it).²⁰ Understood solely as a non-normative project, effective altruism is demanding only insofar as one voluntarily takes on sacrificial projects to maximize the good. These conceptions do not entail that one is morally obligated to do the most good they can, and so are less demanding in this respect.²¹

In *The Most Good You Can Do*, Singer discusses the lives of a number of actual effective altruists,²² all of whom keep some "modest level of comfort and convenience" even if it's

¹⁸ One charitable way of understanding this last criticism is that effective altruism is too close to ideal theory when it should be focused on non-ideal theory.

¹⁹ In a personal correspondence, when presented with each view, Peter Singer endorsed actualism.

²⁰ See Sachs' contribution to this volume for a discussion of whether confronting people with arguments for demanding moral requirements would be counterproductive. Sachs argues that there is "no solid basis" for believing that it would be.

²¹ In his contribution to this volume, MacAskill cites a number of effective altruist surveys, which reveal that 65.4 percent of respondents accept utilitarianism or some other form of consequentialism, and so a substantial percentage of effective altruists likely endorse the normative thesis that one is obligated to do the most good they can. However, 70 percent thought that the definition of effective altruism should be non-normative, and so endorse the non-normative project conception for the purposes of the movement.

²² None of these supposedly exemplary effective altruists live up to the demands of a possibilist effective altruism. But, from Singer's descriptions, they seem to generally live up to the demands of an actualist effective altruism, which is to say that they make choices that will result in them doing the most good *holding fixed the facts about how*

possible for them to do more good by giving up these luxuries.²³ Singer appears to endorse this strategy because choosing to live without this level of comfort and convenience is likely to be “counterproductive”.²⁴ When discussing the sacrifices one must make to be an effective altruist, Singer adds that “if you find yourself doing something that makes you bitter, it’s time to reconsider”.²⁵ He then rhetorically asks whether it would really be best if you chose to do something that made you bitter. Singer’s suggestion appears to be that you ought to avoid such choices because they would result in you doing less good as an effective altruist than if you made choices that prevented you from being bitter. This general advice is given after considering an effective altruist named Julia who decided to have a child in spite of the fact that this would reduce the amount of overall good Julia could do over the course of her life.²⁶

Singer raises similar considerations later in the text when discussing the demandingness of taking a high paying job one doesn’t find intrinsically valuable in order to earn to give. He recognizes that earning to give is “not for everyone” and cautions against it for people who won’t be enthusiastic about “making profits for their employer” even if doing so is necessary for one to do the most good they can over the course of their life.²⁷ Such considerations are echoed by MacAskill in *Doing Good Better*. In the section on what choices prospective effective altruists should make when picking a career, he advises people to pick careers that make them happy, justifying this prescription by pointing out that “if you’re not happy at work, you’ll be less productive and more likely to burn out, resulting in less impact in the long term”.²⁸ In another paper, MacAskill also considers the possibility that taking a high paying job will result in one becoming “corrupted by one’s colleagues” as a result of having to socialize with people who hold anti-effective altruist values.²⁹ He proceeds to suggest ways to mitigate this, and other, risks. MacAskill nevertheless grants that such considerations need to be taken into account by prospective effective altruists.³⁰ He even factors these considerations into the expected value of the choices one makes, which implicitly assumes actualism over possibilism. Similar considerations are raised throughout the chapter on effective altruist career choices. In the concluding chapter, MacAskill advises his readers to set-up a recurring donation at a charity for actualist reasons.³¹

Singer’s actualist assumptions even predate effective altruism. In *The Life You Can Save*, Singer argues that failing to help family members in need “would be going too much against the

they would act if they make any given choice. This is some defeasible evidence that committed effective altruists are disposed toward accepting actualism.

²³ Singer (2016, p. 28).

²⁴ Singer (2016, p. 29).

²⁵ Singer (2016, p. 29).

²⁶ An exact parallel between Singer’s and MacAskill’s actualist response to effective altruists may be found in Peter Railton’s (1984). In that paper, Railton addresses demandingness worries for consequentialism by implicitly assuming actualism. Along the same lines, Singer and Katarzyna de Lazari-Radek briefly mention, and seem to endorse, Frank Jackson’s decision-theoretic consequentialism (1991), which itself assumes actualism. See Singer and de Lazari-Radek (2014, pp. 326-327). Another, not mutually exclusive, possibility is that Singer’s actualist assumptions are rooted in his commitment to esoteric morality. See his and de Lazari-Radek’s (2010). Singer’s remarks in his (2016b: § 2) supports this hypothesis. See also Skelton (2016, pp. 141-142).

²⁷ Singer (2016, p. 47).

²⁸ MacAskill (2016b, p. 149).

²⁹ MacAskill (2014, p. 281).

³⁰ MacAskill (2014, p. 281).

³¹ MacAskill (2015, p. 197).

grain of human nature.”³² He then suggests that people “take care of their families in an entirely sufficient way on much less than they are now spending” and donate the money they have left over to those living in extreme poverty.³³ Actualist considerations even appear to be at the heart of Singer’s “Realistic Approach” to effective giving,³⁴ which sets standards based on what *would* (not what *could*) result in the highest total giving for most individuals (Singer 2009: 154).³⁵

3.2 *The problem with accepting actualism and being an effective altruist*

Although appeals to actualism may help mitigate demandingness worries, actualism also seems to permit actions seemingly antithetical to the commitments of effective altruism understood as a normative claim and as a non-normative project. For instance, actualism entails that Sally is obligated to frivolously spend money on concert tickets today instead of saving three lives tomorrow, but surely Sally fails to do the most good she can by making such a choice. While it’s true that, with respect to her alternatives today, choosing to <purchase the concert ticket> would result in more good than choosing to <keep the money in her account>, it’s also true that Sally can do more good by choosing to <keep the money in her account> today and then choosing to <donate that money to an effective charity> tomorrow. Actualists and possibilists agree that Sally *can* <keep the money in her account and donate that money to an effective charity>. So, actualists and possibilists agree about which acts agents can, and cannot, perform. Their disagreement concerns the relationship between agent’s free actions and their moral obligations. In this case, they disagree about whether Sally is obligated to <keep the money in her account> given that she would then freely decide <not to donate that money to an effective charity>. Consider another example.

Partying Pete: Pete is contemplating gambling away his millions of dollars over a weekend in Vegas. In doing so, he’ll bring some pleasure to himself and friends. Regardless of his intentions today, if Pete does not choose to spend his money in Vegas this weekend, he’ll later decide to spend it on blood diamonds for himself, although he could, at the later time, decide to donate any money he has to an effective charity.

Actualism entails that Pete is obligated to spend his millions partying in Vegas. But surely Pete does not even come close to doing the most good he can by gambling away his millions. After all, actualists and possibilists agree that Pete *can* forgo a Vegas trip and then donate his money to an effective charity. Yet, actualists deny that Pete has an obligation to forgo gambling away his millions in Vegas. Effective altruists who accept actualism inherit the *Not Demanding Enough* and the *Bad Behavior* objections. This is a problem for those who accept effective altruism as a normative thesis because Pete doesn’t seem to come remotely close to doing the most good he can. Yet, if actualism is true, then Pete would be acting in accordance with the normative thesis. This is a problem for those who accept effective altruism as a non-normative project because Pete doesn’t seem to be acting in accordance with the project of doing the most

³² Singer (2009, p. 40).

³³ Singer (2009, p. 40).

³⁴ See, in particular, chapter 10 of Singer’s (2009).

³⁵ Singer also makes remarks along these lines in his (2017, p. 161). Although Singer’s remarks seem to assume actualism, they may also just be the product of his two-level utilitarian view. Similar remarks are made in his (2011, p. 213).

good he can. Yet, if actualism identifies the relevant options, then Pete gambling away his millions is perfectly in line with the effective altruism project. This strikes me as highly implausible. These considerations provide good reason for effective altruists to reject actualism and yet, in turn, threaten to undermine effective altruists' response to demandingness worries.

3.3 Effective altruists' implicit possibilist assumptions

In response to such worries, effective altruists may be inclined to accept possibilism. In fact, in other contexts, effective altruists even implicitly assume possibilism in response to worries that effective altruism is too permissive. Ethical offsetting is one such example. Ethical offsetting is “the practice of undoing harms caused by one's activities through donations or other acts of altruism”.³⁶ Carbon offsetting is an instance of ethical offsetting. A more unconventional example would be personally killing and skinning animals to make a fur coat for oneself, but then donating to non-human animal charities to “cancel out” the number of non-human animal deaths one causes. Killing animals to make a fur coat, however, seems to be in tension with effective altruism understood either as a normative thesis or as a non-normative project. One can do more good by getting others to give up fur, while also giving up fur themselves. In objecting to ethical offsetting on these grounds, effective altruists have implicitly assumed possibilism.³⁷

Effective altruists have repeatedly argued that people should not simply give to charities that are “close to their heart” because doing so is often radically ineffective.³⁸ MacAskill has forcefully argued that one should not let their affinity with a charity (including concern for particular recipients of the charity) even partly determine one's choices about where to give.³⁹ Doing so would be favoring the less important needs of a group of people over the more important needs of another group on the morally irrelevant basis that you happen to know members of one group. Criticisms of the ineffectiveness of using a “close to the heart” heuristic strike me as apt, yet they do implicitly assume possibilism. After all, giving to charities that are “close to the heart” may do more to ensure that one remains sufficiently motivated to give in the future, resulting in them doing more good in the long run than if they now chose to give to an effective charity only to later succumb to *akrasia*.⁴⁰

While anyone working on the ethics of charity or ethical offsetting will have to take a stance on the actualism/possibilism debate, these issues are of special concern for effective altruists. This is because these issues highlight the dilemma effective altruists face.

³⁶ Zabel (2016).

³⁷ See, for example, (Zabel 2016, § 2). Interestingly, while possibilism seems to be implicitly assumed in section 2, actualist considerations are appealed to in the “Caveats” section”. This suggests that both considerations resonate among prospective effective altruists and provides additional reason to accept hybridism.

³⁸ Singer (2009, ch. 4-7); Singer (2016, § 4); MacAskill (2016, ch. 3-5).

³⁹ MacAskill (2015, pp. 41-42). See also (Pummer 2016) for a compelling argument that if one is going to choose to give to charity, they are obligated to give to one that would do the most good even when it would be permissible for them to have not donated at all.

⁴⁰ For a discussion of charities retention strategies and their effects on donor retention rates, see Singer (2009, ch. 4) and Singer (2017, pp. 175-178). See also Schervish and Havens (1997), especially their discussion of the *frameworks of consciousness* motivating factors (p. 241). See also Green and Webb (2008) and Sargeant and Woodliffe (2008). This issue has also been discussed in detail on effective altruist blogs. See, for instance, http://lesswrong.com/lw/6z/purchase_fuzzies_and_utilons_separately/

3.4 The Problem with accepting possibilism and being an effective altruist

Accepting possibilism would have some radical implications for effective altruism. Possibilism combined with the conception of effective altruism as a normative thesis entails that one's present obligations require them to perform the acts that would let them do the most good they can over the course of their entire life independent of considerations about how they would freely act in the future. This may mean that such effective altruists are obligated to forgo donating anything now in favor of investing all their money in order to donate as much as possible on their deathbed even if, on their deathbed, they would choose to donate nothing. It also renders irrelevant the seemingly practical considerations MacAskill and Singer raise about whether one should decide to "earn to give".⁴¹ Most notably, considerations about whether one *would* suffer from "burnout" if they take a labor intensive job that leaves little room for a social life and considerations about whether one *would* become "corrupted" by their co-workers and cease to give to charity are irrelevant to how one is obligated to act if possibilism is true. More generally, any effective altruist who accepts possibilism is committed to the claim that people should completely ignore their motivational structure when determining how to act, sometimes resulting in one foreseeably bringing about the worst possible outcome. This seems antithetical to effective altruism understood as a normative thesis because it requires people to act in ways they know will not only result in them acting wrongly, but also result in the least (not the most) amount of good. This seems antithetical to effective altruism understood as a non-normative project because someone committed to doing to most good shouldn't make choices they know will result in the least amount of good.

Both actualism and possibilism generate commitments seemingly antithetical to each type of effective altruism. Moreover, combining actualism with effective altruism (understood as a normative thesis or as a non-normative project) undermines typical effective altruist responses to ethical offsetting and to donating to charities close to the heart. On the other hand, combining possibilism with effective altruism (understood as a normative thesis) undermines effective altruists' standard response to demandingness worries. Effective altruists are thus faced with a dilemma. I will now argue that the best way out of the dilemma is to accept hybridism.

4. Hybridism

Hybridism retains the benefits of both actualism and possibilism, yet is immune from each of the aforementioned problems. Unlike actualism and possibilism, it's also consistent with the spirit of each type of effective altruism. Recall that hybridism posits a possibilist moral obligation. So, it avoids the *Bad Behavior* and *Not Demanding Enough* objections because it requires agents, such as *Selfish Sally* and *Partying Pete*, to do the best they can throughout their life. At the same time, hybridism avoids the *Worst Outcome Objection* because it posits an actualist practical ought. This component of hybridism is uniquely important for effective altruism understood as a normative thesis because it can address demandingness worries at the practical level.⁴² Hybridism allows that the considerations raised by Singer and MacAskill discussed in section three *are* considerations that one should take into account when deciding how to act. Although,

⁴¹ MacAskill (2015, ch. 9); Singer (2016, ch. 4).

⁴² It's possible that MacAskill and Singer are primarily concerned about addressing demandingness worries at the practical level anyway.

crucially, they are considerations that do not affect one's obligations, but rather, the practical choices one should make in light of their own moral shortcomings.

The ramifications of a hybridism for issues of special concern to effective altruists are worth reviewing. Someone prone to *akrasia* may have extra reason to donate to charities that utilize effective marketing tactics, have high donor retention rates, and are close to one's heart. Contrary to the standard effective altruist suggestion then, people often practically ought to give to suboptimal charities. Similarly, ethical offsetting may frequently be what one practically ought to do, even if it's immoral. One ought to ethically offset (or donate to suboptimal charities) when, of all the acts one can presently ensure she performs, ethical offsetting (or giving to suboptimal charities) would result in the least amount of badness. Adopting hybridism should lead to a more nuanced understanding, and perhaps greater acceptance of, these practices in the effective altruism community.

5. Conclusion

This paper served two goals. First, I explored the implications of the different viable answers to the actualist/possibilist debate for effective altruism. Interestingly, prominent effective altruists seem disposed to endorse the most popular answer (i.e. actualism), which entails that effective altruism is much less demanding than some of its critics believe.⁴³ Yet, there are good reasons to believe that actualism is both false and antithetical to effective altruism, understood either as a normative thesis or a non-normative project. Rejecting actualism, however, undermines common responses effective altruists give to demandingness objections. To make matters worse, there are good reasons to believe that possibilism is also false and antithetical to effective altruism, understood either as a normative thesis or a non-normative project. Rejecting possibilism, however, undermines common arguments effective altruists make against moral offsetting and giving to charities close to the heart. This is the dilemma effective altruists face. Second, I sketched my own positive solution on behalf of effective altruism. Effective altruists can escape this dilemma by adopting hybridism, or so I've argued.⁴⁴

⁴³ See, e.g. Nakul Krishna's (2016), Amia Srinivasan's (2015), and Colin McGinn's (1999). Jeff McMahan reviews, and rebuts, a number of criticisms of effective altruism in his (2016).

⁴⁴ For helpful comments and discussion, I would like to thank Christian Barry, Amy Berg, Donald Bruckner, Alexander Dietz, Karl Ekdahl, Dan Fogal, Johann Frick, Jens Johansson, Michelle Hutchinson, Eden Lin, Theron Pummer, Ketan Ramakrishnan, Olle Risberg, Simon Rosenqvist, Ben Sachs, Peter Singer, Preston Werner, and audiences at the University of St. Andrews, Uppsala University, and the 2018 Philosophy, Politics, and Economics Conference. I owe a special thanks to Hilary Greaves who provided incredibly detailed and insightful written comments on multiple drafts of this paper, which significantly improved the final product. Work on this paper was supported by a grant from the National Endowment for the Humanities and from Seton Hall University's URC summer grant.

References

- Baker, Derek. 2012. "Knowing Yourself – and Giving Up on Your Own Agency in the Process." *Australasian Journal of Philosophy* 90 (4):641-56.
- Bykvist, Krister. 2002. "Alternative Actions and the Spirit of Consequentialism." *Philosophical Studies*, 107 (1): 45-68.
- Cohen, Yishai, and Travis Timmerman. 2016. "Actualism Has Control Issues." *Journal of Ethics and Social Philosophy* 10 (3): 1-19.
- Goldman, Holly S. 1976. "Dated Rightness and Moral Imperfection." *The Philosophical Review*, 85 (4): 449-87.
- Goldman, Holly S. 1978. "Doing the Best One Can." In *Value and Morals: Essays in Honor of William Frankena, Charles Stevenson, and Richard Brandt* by Alvin Goldman. Dordrecht: D. Reidel.
- Greenspan, P.S. 1978. "Oughts and Determinism: A Response to Goldman." *The Philosophical Review* 87 (1): 77-83.
- Green, Corliss, and Deborah Webb. 2008. "Factors Influencing Monetary Donations to Charitable Organizations." *Journal of Nonprofit and Public Sector Marketing* 5 (3): 19-40.
- Jackson, Frank, and Robert Pargetter. 1986. "Ought, Options, and Actualism." *The Philosophical Review* 95 (2): 233-255.
- Jackson, Frank. 1991. "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection." *Ethics* 101 (3): 461-482.
- Jackson, Frank. 2014. "Procrastinate Revisited." *Pacific Philosophical Quarterly* 95 (4): 634-647.
- Karnofsky, Holden. 13 August 2013. "Effective Altruism." *The GiveWell Blog*. (Accessed 11 June 2018). Available at: <https://blog.givewell.org/2013/08/13/effective-altruism/>
- Kiesewetter, Benjamin. 2015. "Instrumental Normativity: In Defense of the Transmission Principle." *Ethics* 125 (4): 921-946.
- Krishna, Nakul. 2016. "Add Your Own Egg." *Examined Life*. (Accessed 23 June 2017). Available at: <http://thepointmag.com/2016/examined-life/add-your-own-egg>
- MacAskill, William. 2015. *Doing Good Better*. New York: Penguin Random House.
- MacAskill, William. 2016. "Banking: The Ethical Career Choice. In *Philosophers Take on the World* edited by D. Edmonds. New York: Oxford University Press.

- MacAskill, William. Forthcoming. "The Definition of Effective Altruism." In *Effective Altruism: Philosophical Issues* edited by Hilary Greaves and Theron Pummer. New York: Oxford University Press.
- McGinn, Colin. 1999. "Our Duties to Animals and the Poor." In *Singer and His Critics* edited by D. Jamieson. Oxford: Blackwell.
- McMahan, Jeff. 2016. "Philosophical Critiques of Effective Altruism." *The Philosopher's Magazine* 73: 92-99.
- Portmore, Douglas. 2011. *Commonsense Consequentialism: Wherein Morality Meets Rationality*. New York: Oxford University Press.
- Portmore, Douglas. Forthcoming. *Opting for the Best*. New York: Oxford University Press.
- Pummer, Theron. 2016. "Whether and Where to Give." *Philosophy and Public Affairs* 44 (1): 77-95.
- Railton, Peter. 1984. "Alienation, Consequentialism, and the Demands of Morality." *Philosophy and Public Affairs* 13 (2): 134-171.
- Ross, Jacob. 2012. "Actualism, Possibilism, and Beyond." In *Oxford Studies in Normative Ethics* by Mark Timmons. New York: Oxford University Press.
- Sachs, Ben. Forthcoming. "Demanding the Demanding." In *Effective Altruism: Philosophical Issues* edited by Hilary Greaves and Theron Pummer. New York: Oxford University Press.
- Sargeant, Adrian, and Lucy Woodliffe. 2008. "Building Donor Loyalty: The Antecedents and Role of Commitment in the Context of Charity Giving." *Journal of Nonprofit and Public Sector Marketing*, 18: 47-68.
- Singer, Peter. 1972. "Famine, Affluence and Morality." *Philosophy and Public Affairs* 1 (1): 229-243.
- Singer, Peter. 2009. *The Life You Can Save*. New York: Random House.
- Singer, Peter, and Katarzyna de Lazari-Radek. 2010. "Secrecy in Consequentialism: A Defence of Esoteric Morality." *Ratio* 23 (1): 34-58.
- Singer, Peter. 2011. *Practical Ethics*. New York: Cambridge University Press.
- Singer, Peter, and Katarzyna de Lazari-Radek. 2014. *The Point of View of the Universe: Sidgwick and Contemporary Ethics*. New York: Oxford University Press.
- Singer, Peter. 2015. *The Most Good You Can Do*. New York: 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.
- Singer, Peter. 2017. *Ethics in the Real World*. New York: Princeton University Press.
- Skelton, Anthony. 2016. "The Ethical Principles of Effective Altruism." *Journal of Global Ethics* 12 (2): 137-146.
- Schervish, Paul G., and John J. 1997. "Social Participation and Charitable Giving: A Multivariate Analysis." *Voluntas*, 8 (3): 235-260.
- Sobel, Jordan Howard. 1976. "Utilitarianism and Past and Future Mistakes." *Noûs* 10 (2): 195-219.
- Srinivasan, Amia. 2015. "Stop the Robot Apocalypse." *London Review of Books* 37 (18): 3-6.
- Timmerman, Travis. 2015a. "Sometimes There is Nothing Wrong with Letting a Child Drown." *Analysis* 75 (2): 204–212.
- Timmerman, Travis. 2015b. "Does Scrupulous Securitism Stand-Up to Scrutiny? Two Problems for Moral Securitism and How We Might Fix Them." *Philosophical Studies* 172 (6): 1,509-1528.
- Timmerman, Travis, and Yishai Cohen. 2016. "Moral Obligations: Actualist, Possibilist, or Hybridist?" *Australasian Journal of Philosophy* 94 (4): 672-686.
- Timmerman, Travis, and Philip Swenson. Forthcoming. "How to be an Actualist and Blame People." In *Oxford Studies in Agency and Responsibility, Vol. 5* by David Shoemaker. New York: Oxford University Press.
- Wedgwood, Ralph. 11 September 2009. "Against Actualism." *PEA Soup*. (Accessed 20 May 2017). Available at: peasoup.typepad.com/peasoup/2009/09/against-actualism.html
- Woodard, Christopher. "What's Wrong with Possibilism." *Analysis* 69 (2): 219-26.
- Zabel, Claire. 5 January 2016. "Ethical Offsetting is Antithetical to Effective Altruism." *Effective Altruism Forum*. (Accessed 20 May 2017). Available at: http://effective-altruism.com/ea/ry/ethical_offsetting_is_antithetical_to_ea/
- Zimmerman, Michael. 1996. *The Concept of Moral Obligation*. New York: Cambridge University Press.
- Zimmerman, Michael. 2006. "The Relevant Risks to Wrongdoing." In *The Good, the Right, Life and Death: Essays in Honor of Fred Feldman* by Kris McDaniel, Jason Raibley, Richard Feldman, and Mark Zimmerman. Burlington: Ashgate Publishing Co
- Zimmerman, Michael. 2017. "Prospective Possibilism." *The Journal of Ethics* 21 (2): 117-150.