

Supererogation, Suberogation, and Maximising Expected Choiceworthiness

(This is a pre-print version. Please cite final version forthcoming in *Canadian Journal of Philosophy*.)

Leora Urim Sung

Philosophy Department, University College London, UK

leora.sung.17@ucl.ac.uk

Abstract

Recently, several philosophers have argued that, when faced with moral uncertainty, we ought to choose the option with the maximal expected choiceworthiness (MEC). This view has been challenged on the grounds that it is implausibly demanding. In response, those who endorse MEC have argued that we should take into account the *all-things-considered* choiceworthiness of our options when determining the maximally choiceworthy option. In this paper, I argue that this gives rise to another problem: for the most part, acts that we consider to be supererogatory are rendered impermissible, and acts that we consider to be suberogatory are rendered obligatory, under MEC. This problem arises because, when we factor in prudential reasons, we often have most reason, or most expected reason, to act in accordance with our interests. I suggest a way to reformulate MEC so that prudential reasons only make acts permissible or non-obligatory, without ever making acts obligatory or wrong under moral uncertainty.

Keywords: moral uncertainty, maximising expected choiceworthiness, supererogation, suberogation

1. Introduction

Being imperfect beings, we often need to make decisions under uncertainty about a vast range of facts. These include not only descriptive facts, but also moral facts. When we are uncertain about descriptive facts in a moral setting, we can distinguish between what we *objectively* ought to do and what we *subjectively* ought to do.¹ What we objectively ought to do is whatever

¹ See Parfit (2011).

morality would require of us if we were aware of all the relevant descriptive facts. What we subjectively ought to do is what morality requires of us given our descriptive uncertainty. *Moral Uncertainty* is the idea that there is an ‘ought’ which speaks to what we should do when we face a moral decision and we are uncertain about which moral principles are true. This is sometimes referred to as what we *super-subjectively* ought to do.² In Section Two, I say more about the nature of this super-subjective ought.

Maximising Expected Choiceworthiness (MEC) is a theory of decision-making under moral uncertainty, and philosophers who endorse this position in some form include William MacAskill (2018, 2020), Krister Bykvist (2020), Toby Ord (2020), Jacob Ross (2006), and Andrew Sepielli (2009). However, MEC faces many challenges. One important challenge, which I explain below and then set aside, is the *problem of intertheoretic value comparisons*. But MEC has also been challenged on the grounds that it is implausibly demanding. Those who endorse MEC have defended the theory against this worry by making it sensitive to prudential reasons, introducing what they call the *all-things-considered* choiceworthiness ordering of options. I explain this in Section Three.

In Section Four, I argue that, while taking into account prudential reasons lessens the demandingness of MEC, it makes MEC vulnerable to another objection: for the most part, when we consider prudential reasons in determining the choiceworthiness of options, acts which we consider to be supererogatory are rendered super-subjectively impermissible, and acts which we generally consider to be suberogatory are rendered super-subjectively obligatory.

In Section Five, I examine a potential solution to the problem raised in Section Four. This solution holds that the super-subjective ‘ought’—the ‘ought’ of moral uncertainty—refers only to what we purely rationally ought to do rather than to what we morally ought to do. I argue that this response should be rejected for several reasons.

In Section Six, I suggest a way of formulating MEC that allows us to choose how much weight to give to our prudential reasons. I show that this formulation, which I call *Discretionary* MEC, is able to provide the correct intuitive results regarding the permissibility of supererogatory acts and the optionality of suberogatory acts.

In Section Seven, I explain that Discretionary MEC faces a limitation when making room for all supererogatory acts, and I suggest some ways in which we can potentially overcome this limitation.

² For a clear explanation of different levels of moral ‘ought’, see Hedden (2016).

2. The ‘Ought’ of Moral Uncertainty

It seems plausible that there is an ‘ought’ which speaks to what we should do under moral uncertainty.³ Consider the following case:

Drug. A and B suffer from a fatal illness. A is a human patient, whereas B is a chimpanzee. You have a vial of a drug that can help. If you administer all of the drug to A, she will be completely cured. If instead you split the drug equally between A and B, both will be completely cured. You are uncertain about the moral status of non-human animals and think it is equally likely that non-human animals have no moral value and that non-human animals have the same moral value as humans.

What should you do? It seems that, regardless of whether non-human animals have moral value, there is a sense in which you ought to split the drug equally between A and B. This is because splitting the drug is the more choiceworthy option according to moral theories which grant moral value to non-human animals, and it is at least as choiceworthy according to those moral theories which do not. Given your moral uncertainty, it would be inappropriate to administer all the drug to A, knowing that there is a risk of doing something that is morally impermissible, when, instead, you can choose another option which guarantees you do no wrong.

The ‘ought’ of moral uncertainty is clearly different from the ‘ought’ of first-order moral theories. Perhaps it is true that non-human animals have no moral value, and so choosing to administer all the drug to A while leaving B to die is morally permissible. But doing so is evidently not what you ought to do given your moral uncertainty.

But if the ‘ought’ of moral uncertainty is not the same as the ‘ought’ of first-order moral theories, what kind of ‘ought’ is it?

One way to make sense of the ‘ought’ of moral uncertainty is to say that there are different levels of moral ‘ought’. As mentioned earlier, we can distinguish between what we *objectively* ought to do when we know all the morally relevant details of the situation, and what we *subjectively* ought to do when we are uncertain about the descriptive facts that are morally relevant. The ‘ought’ of moral uncertainty can be understood as one step further up—as what we *super-subjectively* ought to do when we are uncertain about which moral theory is true. In

³ Some philosophers argue that this intuition is misguided. See, for example, Harman (2014) and Weatherson (2019). I set aside these worries here.

other words, while first-order moral theories answer the question of what we morally ought to do at the first level, the ‘ought’ of moral uncertainty—the super-subjective ‘ought’—speaks to the question of what we morally ought to do when we are not sure what we morally ought to do at the first level.

Another way to make sense of the ‘ought’ of moral uncertainty is in terms of rationality—it is what you rationally ought to do, given your beliefs and preferences. However, it seems that the super-subjective ought cannot refer to a *purely* rational ‘ought’. For example, it can be rational for me to administer all the drug to A under moral uncertainty; it may be that I simply do not care about doing what is morally right, and I prefer at that moment to use all the drug on A. So, if we are to interpret the super-subjective ought in terms of rationality, it only applies to moral agents with certain preferences. The ‘ought’ of moral uncertainty refers to what the morally conscientious person—someone who cares about doing right and refraining from doing wrong—rationally ought to do, given their beliefs and preferences.

So, on this alternative view, the ‘ought’ of moral uncertainty refers to the ought relevant to a morally conscientious person, which comprises of both the moral and the rational. The moral element requires you to be a morally conscientious person, caring about doing right and refraining from doing wrong. The rational element requires you to act in ways that cohere with these beliefs and preferences. So, the rational and morally conscientious person is required to act in ways that mitigate moral risk under conditions of moral uncertainty. If you are morally conscientious and rational, you should split the drug between A and B, because you care about doing the right thing and splitting the drug will guarantee that you do the right thing.

I will follow the latter interpretation of the super-subjective ought in the rest of the paper. However, the arguments that I make will also apply if we drop the rational element and make sense of the super-subjective ought only in terms of morality. The only account of the super-subjective ought that would conflict with my arguments would be the purely rational account. I have just suggested why we should not accept this view, and I will further address this in Section Five.

Now that we have defined what the ‘ought’ of moral uncertainty refers to, how do we find out what the ‘ought’ of moral uncertainty prescribes? In the next section, I will introduce the most prominent theory of decision-making under moral uncertainty—Maximising Expected Choiceworthiness—and outline a problem the theory faces.

3. The Demandingness of MEC

In cases of descriptive uncertainty, many accept that one should choose the option which has the maximal expected utility. Some philosophers are in favour of treating descriptive and moral uncertainty analogously, and they argue that under conditions of moral uncertainty, we should choose the option with the maximal expected choiceworthiness, where choiceworthiness represents the strength of the reasons for choosing an option (MacAskill et al. 2020, 4).

Maximise Expected Choiceworthiness (MEC): when we can determine the expected choiceworthiness of options, we ought to choose the option which has the maximal expected choiceworthiness.⁴

MacAskill et al., define permissibility as optimal choiceworthiness—i.e., an act is permissible if and only if it is maximally choiceworthy. And when we can determine the choiceworthiness of options, an act is *super-subjectively* permissible—that is, permissible under conditions of moral uncertainty—if and only if it has the maximal expected choiceworthiness.

What if a moral theory says that two options are both permissible but that one is *supererogatory*, meaning that this option is, in some way, morally superior to the other? Are both options equally maximally choiceworthy, or is the supererogatory option more choiceworthy than the other? MacAskill et al. suggests that under the most plausible accounts of supererogation, both options are equally choiceworthy. If instead we hold the view that while both options are permissible, the supererogatory act is *more* choiceworthy than the merely permissible act, this results in problems for MEC. For example, imagine you are certain in the moral theory on which both options are morally permissible but one is supererogatory. MEC then implies that this supererogatory act is super-subjectively obligatory because it is more choiceworthy, even though you are certain that this act is not objectively obligatory. This is clearly counterintuitive, and so it seems we should regard all permissible acts to be equally morally choiceworthy when we are considering the choiceworthiness of options.

⁴ MacAskill et al. use the term “appropriate” to make assessments of options under moral uncertainty and define MEC in the following way: when we can determine the expected choiceworthiness of different options, *A* is an appropriate option if and only if *A* has the maximal expected choiceworthiness. *A* is more appropriate than *B* if and only if a rational and morally conscientious agent who had the same set of options and beliefs would prefer *A* to *B*. See MacAskill et al. (2020), p. 48.

One common challenge to MEC comes from the problem of intertheoretic value comparisons.⁵ In order to determine the expected choiceworthiness of a certain moral option, there needs to be some non-arbitrary basis for comparing degrees of moral value or disvalue attributed to this option by competing moral theories. But although it is possible to compare value differences within a theory, it seems impossible to do so across different theories, as there is no common scale shared by both. Much of the literature on the topic of moral uncertainty aims to provide a solution to this problem of intertheoretic value comparisons, but all the proposals face compelling objections.⁶ I will, however, assume that the problem of intertheoretic value comparisons can be solved in some way, in order to focus on a different problem that MEC faces.

MEC has also been criticized for being overly demanding: it has implications that require too great a personal sacrifice from us.⁷ Consider the following case:

Burning Building. Two people are trapped inside a burning building, and there is no one around to help but you. You can enter the building and save two people's lives, but only at the cost of your own life.

There are two possible options available to the moral agent: SELF-SACRIFICE or DO NOTHING. Suppose that you have high credence that both of these possible options are permissible, with SELF-SACRIFICE thus being permissible but not required. Suppose, also, that you have some small credence in act utilitarianism, according to which SELF-SACRIFICE is the only permissible option. MEC then seems to imply that SELF-SACRIFICE is the only super-subjectively permissible option, as this is the option with the maximal expected choiceworthiness. However, this seems intuitively implausible. Most would agree that this conclusion requires too much of the moral agent, and therefore, should be rejected.

It could be responded that the demandingness of MEC is really a problem with demanding theories like act utilitarianism, rather than with MEC itself. According to this response, we should consider the demandingness of theories when deciding what credence we

⁵ For good overviews of the moral uncertainty literature, see Bykvist (2017) and MacAskill et al. (2020).

⁶ For potential solutions to the problem of intertheoretic value comparisons, see Lockhart (2000), Ross (2006), and Sepielli (2009). For some objections, see Gracely (1996), Gustafsson and Torpman (2014), and Hedden (2016).

⁷ See Weatherson (2002); Barry and Tomlin (2016); Rosenthal (2020).

ought to have across different moral theories, but not when we are evaluating MEC. If we reject MEC on the grounds of it being too demanding, it seems we are guilty of ‘double-counting’, because we are allowing our intuitions about demandingness to reduce our credence both in first-order moral theories and also in second-order theories like MEC (MacAskill et al. 2020, 51).

I am not persuaded by this response. We not only have intuitions about what we objectively or subjectively ought to do, but also have intuitions about what we super-subjectively ought to do. When we object to the demandingness of first-order moral theories, we are appealing to the intuition that what we objectively ought to do is less demanding than what these theories imply. When we object to the demandingness of MEC, we are appealing to the intuition that what we super-subjectively ought to do is less demanding than what MEC implies. So, we are relying on different intuitions, rather than double counting. (Consider the following analogy. It does not seem to be double counting to reduce credence in theories of the super-subjective ought that tells me to do what is impossible, even after reducing credence in first-order moral theories that tell me to do the impossible. And it seems that ‘ought implies not-overly-demanding’ is close enough in form and content to ‘ought implies can’ that we can have intuitions about them both as principles governing the super-subjective ought.⁸)

MacAskill, Bykvist, and Ord also offer a second response to the demandingness objection to MEC (2020, 52–53). They argue that an account of decision-making under moral uncertainty should take into account what the *all-things-considered* choiceworthiness ordering is. By this they mean we ought to take into account non-moral reasons, such as prudential reasons, as well as moral reasons, when determining the choiceworthiness of options. For the sake of simplicity, I will refer to the all-things-considered version of MEC as *Expanded MEC*.

If we accept Expanded MEC, MEC may not be unreasonably demanding. For instance, Expanded MEC will not necessarily require you to choose SELF-SACRIFICE, even if you have some credence in act utilitarianism. This is because, presumably, you will have reasonable credence in the view that act utilitarianism is wrong *and* you have non-moral reasons, or prudential reasons, to preserve your own life. If you do, then, what you should do, according to MEC, will depend on how plausible you find this view relative to act utilitarianism. If your credence in act utilitarianism is sufficiently small, and your credence in the view that you have prudential reasons to stay alive is sufficiently high, it would be appropriate for you to choose

⁸ I thank Timothy Luke Williamson for this suggestion.

DO NOTHING instead. The option with the maximal expected choiceworthiness would be for you to refrain from sacrificing your life.

Although broadening the scope of consideration to include prudential reasons allows MEC to escape the demandingness objection, it seems to give rise to another problem. I now turn to explain this problem.

4. A Problem with “All-Things-Considered Choiceworthiness”

In response to the demandingness objection to MEC, MacAskill et al. introduced a version of MEC which takes into account the all-things-considered choiceworthiness ordering of options. This version of MEC, Expanded MEC, includes prudential reasons as well as moral reasons when calculating which option has the maximal expected choiceworthiness. This revision succeeds in lessening the demandingness of MEC. However, I will now argue that when we include prudential reasons in determining the choiceworthiness of options, another problem arises: for the most part, Expanded MEC results in counter-intuitive implications when it comes to supererogatory and suberogatory acts.

4.1. Suboptimal Supererogation

Supererogation is a term for a class of acts which go “beyond the call of duty”. There is not one agreed definition of supererogation, but most agree that supererogatory acts must be *permissible yet optional*, and, in some sense, morally *better* in comparison to other, available morally permissible acts.⁹

If we consider only moral reasons when determining the choiceworthiness of options, MEC would frequently require us to perform acts which we consider to be supererogatory. This is because even if we are almost certain that we are permitted but not required to perform the intuitively supererogatory act, we are likely to also have some credence in moral theories which deem such acts morally obligatory. So, the intuitively supererogatory act will be the option with the maximal expected choiceworthiness, and therefore, super-subjectively required according to MEC. This is the demandingness objection.

⁹ See Urmson (1958).

However, if we instead follow Expanded MEC, by including prudential reasons as well as moral reasons in our calculation of the choiceworthiness of options, then if the prudential reasons against a particular option are sufficiently strong, this would make that option *suboptimal* in terms of expected choiceworthiness. Expanded MEC would then imply that we ought to refrain from choosing this suboptimal option. It follows that certain acts that we consider to be supererogatory at the first-order level can turn out to be super-subjectively *impermissible*, because of our prudential reasons weighing against performing such acts. I call this *the problem of suboptimal supererogation*.

For example, we said that under the original formulation of MEC, without the all-things-considered choiceworthiness ordering, you are super-subjectively required to choose SELF-SACRIFICE—that is, the intuitively supererogatory option of sacrificing your life to save two people. This is the demandingness objection. If we shift to a version of MEC that does take into account prudential considerations, Expanded MEC, we can avoid this objection. However, this new version of MEC goes too far in the other direction. Because it takes prudential reasons into consideration, Expanded MEC will *prohibit* you from choosing SELF-SACRIFICE, and require you to choose DO NOTHING instead. This seems to be intuitively implausible. Although we may not be *obligated* to sacrifice our life, intuitively, it seems we should at least be permitted to do so, even under moral uncertainty.

Suppose that Expanded MEC dictates that you give less weight to your prudential reasons, so that the option with the maximal expected choiceworthiness is SELF-SACRIFICE rather than DO NOTHING. It will then be super-subjectively permissible for you to choose SELF-SACRIFICE. However, Expanded MEC now fails to avoid the demandingness objection, as you are not only super-subjectively *permitted*, but also super-subjectively *required* to choose SELF-SACRIFICE. So, there is a dilemma here for Expanded MEC. It must dictate that our prudential reasons be weighty enough to avoid the demandingness objection. But if our prudential reasons are weighty enough to avoid the demandingness objection, they will also weigh against performing intuitively supererogatory acts, prohibiting us from performing them under moral uncertainty.

To restate the problem, although the all-things-considered choiceworthiness ordering rescues MEC from the demandingness objection, it seems to go too far, by not only exempting us but also *prohibiting* us from actions which we consider to be supererogatory.

4.2. Optimal Suberogation

Expanded MEC faces a similar problem when it comes to suberogation. Suberogatory acts are a mirror opposite of supererogatory acts; they must be *permissible yet optional*, and, in some sense, morally *worse* than other permissible acts. Julia Driver (1993), who first coined the term, defines suberogation as acts which are bad to do, but not forbidden.

If we consider only moral reasons when determining the choiceworthiness of options, MEC would frequently prohibit us from performing acts which we consider to be suberogatory. This is because even if we are almost certain that we are permitted but not required to do the intuitively suberogatory act, it is likely that we will have at least some credence in moral theories which deem such acts morally impermissible. So, the intuitively suberogatory act will not be the option with the maximal expected choiceworthiness, and it will thus be super-subjectively impermissible according to MEC.

However, if we instead follow Expanded MEC, by including prudential reasons as well as moral reasons in our calculation of the choiceworthiness of options, then if the prudential reasons in favour of a particular option are sufficiently strong, this would make that option *optimal* in terms of expected choiceworthiness. Expanded MEC would then imply that we are required to choose this optimal option. This means that certain acts that we consider to be suberogatory on the first-order level of morality would be considered super-subjectively *obligatory*, because of our prudential reasons weighing in favour of performing such acts. I call this *the problem of optimal suberogation*.

For example, say you have a sibling who is dying of kidney failure, and you are the only person with a compatible kidney. If we take it that the act of refusing to donate your kidney is morally permitted, but in some way, morally worse than donating your kidney, it will fall under the category of the suberogatory. However, it is difficult to see how MEC could accommodate this. Suppose you are almost certain that you are not morally obligated to donate your kidney, and you also have strong prudential reasons not to do so. If you do not factor in prudential reasons into your calculation of expected choiceworthiness, the option of donating your kidney will have the maximal expected choiceworthiness, and so be super-subjectively obligatory. But if you *do* factor in prudential reasons, then, given the personal costs involved and your low credence in the view that there is a moral obligation to donate your kidney, leaving your sibling to die will be the option which has the maximal expected choiceworthiness. According to Expanded MEC, then, it is super-subjectively *impermissible* for you to donate your kidney because it is not the option with the maximal expected choiceworthiness. Rather, the option with the all-things-considered maximal expected choiceworthiness would be to leave

your sibling to die, and so you are super-subjectively required to do what most would regard to be supererogatory.

4.3. Generalising the Problem

To sum up the problem, on the one hand, Expanded MEC renders super-subjectively impermissible acts which the moral agent is almost certain are supererogatory, because these acts are all-things-considered suboptimal in terms of expected choiceworthiness. When this is the case, Expanded MEC prohibits the agent from doing something which the moral agent is almost certain is morally good. On the other hand, acts which the moral agent considers to be supererogatory are rendered super-subjectively obligatory under Expanded MEC, because these acts are all-things-considered optimal in terms of expected choiceworthiness. Here, Expanded MEC requires the agent to do something the moral agent is almost certain is morally bad.

This problem arises because of two features of Expanded MEC. First, like all maximising theories, Expanded MEC conflates obligation and permissibility in most cases. Unless there are multiple options with the maximal expected choiceworthiness, whatever is permissible is also obligatory. Supererogatory and supererogatory acts need to be permissible yet optional, but Expanded MEC does not allow for such a category of moral acts unless there are ‘ties at the top’, with more than one option with the maximal expected choiceworthiness. The second feature of Expanded MEC that gives rise to the problem is that your prudential reasons are factored into determining the choiceworthiness of options. As prudential reasons play a role in making an option maximally choiceworthy, and as we are obligated to choose the option with the maximal expected choiceworthiness, the result is that we are often obligated to act in accordance with our interests and prohibited from acting against our interests.¹⁰ It is because of this feature of Expanded MEC that supererogatory acts are rendered super-subjectively impermissible, and supererogatory acts are rendered super-subjectively obligatory.

However, it seems that while our prudential reasons may excuse us for not performing the morally best act, they should not prohibit us from performing them, i.e., make them impermissible. It makes sense to say that if you have strong prudential reasons against option *A*, you are not *obligated* to do *A*, but we should not go so far as to say that you are not *permitted* to do *A*. Similarly, prudential reasons may make a certain act permissible that would have been

¹⁰ This is also something of a failing of traditional utilitarian theories—they entail that prudence is morally required when your own welfare is the only thing that is at stake.

impermissible all-things-being-equal, but they should not go so far as to make the act obligatory. If you have strong prudential reasons to choose *A*, although you may be permitted to do *A*, you should not be *obligated* to choose *A*.

This is because we ordinarily think that there is an asymmetry with regards to what an agent is permitted or obligated to do to herself and what she is permitted or obligated to do to others.¹¹ While other-regarding reasons may prohibit you or require you to perform a certain act, it seems that self-regarding reasons should not.¹² With the all-things-considered choiceworthiness ordering, however, self-regarding reasons affect what the option with the maximal expected choiceworthiness is, so that we are prohibited from acting against our interests to a certain degree and obligated to act in accordance with our interests to a certain degree. This makes selfless acts impermissible and self-interested acts obligatory under moral uncertainty.

In Section Six, I offer a solution to the problems I have raised for MEC. But before doing so, in the next section, I will consider an alternative potential solution.

5. A Potential Solution: Rational Ought, Not Moral Ought

So far, I have argued that Expanded MEC produces counter-intuitive implications when it comes to cases of supererogation and suberogation. This is because acts which we generally consider to be supererogatory are rendered suboptimal when we factor in our prudential reasons, making such acts super-subjectively impermissible. Correspondingly, acts which we generally consider to be suberogatory are rendered optimal when we take our prudential reasons into account, making such acts super-subjectively obligatory.

Could we solve the problem of suboptimal supererogation and optimal suberogation by re-interpreting the ‘ought’ of moral uncertainty as a purely *rational* ‘ought’ rather than the ‘ought’ that is relevant to the rational and morally conscientious agent? If so, when we say that an act is super-subjectively impermissible, we are saying that the act is merely rationally impermissible, rather than impermissible in any moral sense. When we say that an act is super-subjectively obligatory, we are saying that an act is simply required by rationality, rather than it being obligatory in any moral sense.

¹¹ For discussions on this self-other asymmetry, see Stocker (1976) and Slote (1984).

¹² It is nonetheless still somewhat controversial whether there are duties to oneself. For a defence of self-regarding duties, see Muñoz and Baron-Schmitt (2024).

This interpretation solves the problem of suboptimal supererogation in the following way: supererogatory acts are morally permissible and good, but when such acts are the suboptimal act under Expanded MEC, these acts are rationally impermissible. Although it is morally good for the moral agent to choose SELF-SACRIFICE, given her prudential reasons to stay alive and her low credence in any moral theory on which SELF-SACRIFICE is morally required, it is not what she rationally ought to do under moral uncertainty. What the agent rationally ought to do would be to choose DO NOTHING, as this is what the all-things-considered choiceworthiness version of MEC prescribes. This distinction between moral and rational ‘ought’s allows us to maintain the *moral* permissibility of choosing options which the moral agent regards as supererogatory, while also saying that MEC would dictate that you ought not to choose such options *rationally* speaking.

The same response can be given to the problem of optimal suberogation. We can say that suberogatory acts are morally worse than other permissible acts, but if it is the optimal option under Expanded MEC, it is rationally required of the agent. Although it is morally worse for you to refuse to donate your kidney, it is what you rationally ought to do, given your prudential reasons and given your low credence that it is morally required. This allows us to maintain that the act of refusing to donate your kidney is *morally* worse, while also saying that it is *rationally* required by MEC. Although suberogatory acts are worse in that the agent could have done better, morally speaking, it is rationally obligatory for the agent to choose the suberogatory option under moral uncertainty.

This kind of response should be rejected for several reasons. First, it would undermine the need to take moral uncertainty seriously in ethics. If the ‘ought’ of moral uncertainty refers only to what it is rational to do, rather than what is rational for a morally conscientious agent to do, there seems to be no moral force behind MEC. MEC just becomes a theory of what we rationally ought to do given our beliefs and preferences, no longer having anything to do with mitigating moral risk—unless we happen to care about doing so. The super-subjective ought, when interpreted in terms of rationality, cannot refer to a purely rational ought, but a rational ought relevant to a morally conscientious person. Also, in cases where the demands of rationality clash with the demands of morality, why should we follow the demands of rationality over those of morality? When making moral decisions, it seems our concern, one way or another, should be with moral rightness. So, if we understand the super-subjective ought in purely rational terms, it seems we should just ignore what MEC tells us about what it is rational to do and follow what we believe is required by morality. Finally, this solution contradicts the position endorsed by MacAskill et al. They claim that the agent who is both

rational *and* morally conscientious would maximise expected choiceworthiness. This is because the morally conscientious person would both care about reducing moral risk and not act in a way that is morally risk-taking. In other words, it is the morally conscientious person who cares about doing the right thing that would act in accordance with MEC in order to reduce moral risk.

So, if we accept this potential solution, we are left with counterintuitive portrayals of moral agents. The agent who goes against her prudential reasons to undertake a significant personal sacrifice in order to do what is morally good would be acting contrary to what the morally conscientious person rationally ought to do. The agent who gives a large proportion of her income to charity, for example, is morally praiseworthy, yet irrational or morally unconscientious. Refusing to donate your kidney to save your dying sibling is morally reproachable, yet you ought to refuse, because this is what the rational and morally conscientious person would do.

In the next section, I develop a different solution to the problem of suboptimal supererogation and the problem of optimal suberogation. I believe this solution gets to the root of the problem by ensuring that prudential reasons only make acts super-subjectively permissible or non-obligatory, without ever making acts super-subjectively obligatory or wrong.

6. My Solution: Discretionary MEC

As I argued above, the critical problem with Expanded MEC is that your prudential reasons are factored into determining the choiceworthiness of options, so that you are often prohibited from acting against your interests and obligated to act in accordance with your interests. This renders many selfless acts super-subjectively impermissible and many self-regarding acts super-subjectively obligatory.

When it comes to the first-order level of morality, it is generally accepted that while your prudential reasons may exempt you from being morally required to choose an option *A*, they should not go so far as to prevent you from choosing *A*. Similarly, while your prudential reasons may permit you to choose an option *A*, they should not go so far as to obligate you to choose *A*. It thus seems plausible that, under moral uncertainty, while your prudential reasons may exempt you from being *super-subjectively* required to choose an option *A*, they should not go so far as to make it *super-subjectively* impermissible for you to choose *A*. And while your

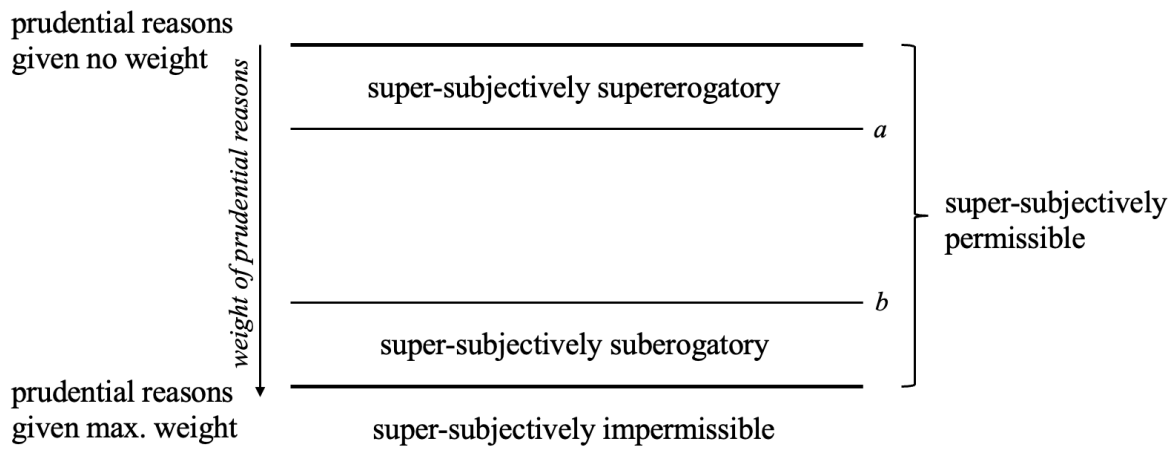
prudential reasons may make it *super-subjectively* permissible to choose option *A*, they should not make it so that it is *super-subjectively* required for you to choose *A*.

There is a natural and interesting way of reframing MEC so that prudential reasons only exempt you from performing, or give you permission not to perform, an act, while not preventing you from performing, or obligating you to perform, the act. On Expanded MEC, prudential reasons are given the maximum permissible weight, whereas on the original version of MEC, prudential reasons are given no weight at all. I propose a new version of MEC, on which the agent can *choose* how much weight to give to her prudential reasons. By this, I mean it is up to the agent whether to give maximum weight to her prudential reasons (by which I mean whatever weight Expanded MEC permits the agent to give to her prudential reasons), or to give less weight to her prudential reasons, or to exclude her prudential reasons entirely when determining the choiceworthiness of her options. As this view leaves it up to the agent how much weight to give to her prudential reasons, I call this view *Discretionary* MEC.

Under Discretionary MEC, there can be many options which have the maximal expected choiceworthiness, because the agent can choose to give varying amounts of weight to her prudential reasons. We can say that acts are super-subjectively *permissible* if and only if they have the maximal expected choiceworthiness when the agent gives her prudential reasons their full weight, or when the agent gives her prudential reasons some of their weight, or when the agent gives her prudential reasons no weight at all. Conversely, we can say that an act is super-subjectively *impermissible* if and only if it does not have the maximal expected choiceworthiness under Discretionary MEC—that is, if and only if it is the suboptimal option regardless of whether the agent gives full, partial, or no weight to her prudential reasons. This means that whatever option has the maximal expected choiceworthiness when prudential reasons are given maximum weight would be the very least that is required of the moral agent. As the agent is not permitted to give her interests more weight than what is maximally permissible, anything less than what is required of the agent when prudential reasons are given maximum weight would be super-subjectively impermissible.

How does this reformulation solve the problems I raised for MEC and Expanded MEC? As I said, according to Discretionary MEC, it is up to the agent how much weight to give to her prudential reasons. So, in many cases, there will be multiple options with the maximal expected choiceworthiness depending on the weight she assigns to her prudential reasons. These options are all super-subjectively permissible because they all have the maximal expected choiceworthiness given some permissible weighting of prudential reasons. So, Discretionary MEC can accommodate our intuitions about supererogatory and suberogatory

acts by making room for *super-subjectively* supererogatory acts and *super-subjectively* suberogatory acts—that is, acts which are supererogatory and suberogatory under conditions of moral uncertainty. We can say that an act is super-subjectively *supererogatory* if and only if it has the maximal expected choiceworthiness only when the agent gives her interests significantly less than their full weight. And we can say that an act is super-subjectively *suberogatory* if and only if it has the maximal expected choiceworthiness only when the agent gives her interests close to their maximum weight. The following diagram provides an illustration of the different categories of acts under Discretionary MEC:



The scale on the left shows the varying degrees of weight the agent can give to her prudential reasons. At the top, prudential reasons are excluded entirely, with the weight we give to our prudential reasons gradually increasing down to the bottom, where prudential reasons are given their maximum permissible weight. On the right, we have different acts which have the maximal expected choiceworthiness depending on how much weight the agent chooses to give to her prudential reasons. If we accept Discretionary MEC, so long as an act has the maximal expected choiceworthiness when the agent gives her prudential reasons their full weight, partial weight or no weight, the act is super-subjectively permissible. An act is super-subjectively *supererogatory* if and only if it has the maximal expected choiceworthiness only when the weight assigned to her prudential reasons falls below a certain level, *a*, down until where prudential reasons are given no weight at all. Correspondingly, an act is super-subjectively *suberogatory* if and only if it has the maximal expected choiceworthiness only when the weight assigned to her prudential reasons is above a certain level, *b*, up until where prudential reasons are given their maximum permissible weight. Acts which have the maximal expected choiceworthiness only when agents give *greater* weight to their prudential reasons than the maximum permissible weight would be super-subjectively *impermissible*.

Not only does this reformulation of MEC allow us to maintain the permissibility and optionality of supererogatory acts, it also allows us to satisfy all the conditions of supererogation, but at the level of moral uncertainty. For an act to be supererogatory, we said it needs to be permissible yet optional, and in some sense, morally better than other morally permissible acts. Under Discretionary MEC, super-subjectively supererogatory acts are super-subjectively permissible because they have the maximal expected choiceworthiness when the agent gives her own interests significantly less weight. Such acts are also optional, because the agent can instead choose to assign different weight to her interests, resulting in a different option with the maximal expected choiceworthiness. Finally, super-subjectively supererogatory acts are, in some sense, morally better than options which have the maximal expected choiceworthiness only when the agent gives greater weight to her prudential reasons—they are better in terms of other-regarding reasons.

The same goes for suberogation. For an act to be suberogatory, we said it needs to be permissible yet optional and in some sense, morally worse than other morally permissible options. Under Discretionary MEC, super-subjectively suberogatory acts are super-subjectively permissible because they have the maximal expected choiceworthiness when we choose to grant significant weight to our prudential reasons. As the agent can assign different weight to her interests, resulting in other options which have the maximal expected choiceworthiness, such suberogatory acts are also optional. Finally, there is some sense in which these super-subjectively suberogatory acts are also morally worse than other permissible acts, as they are worse in terms of other-regarding reasons than options which have the maximal expected choiceworthiness when the agent assigns less weight to her prudential reasons.

Reformulating MEC in this way solves the problem of suboptimal supererogation and the problem of optimal suberogation. Discretionary MEC permits a moral agent to sacrifice her life to save two people, for example, even though this is not the option with the maximal expected choiceworthiness under Expanded MEC. This is because this intuitively supererogatory act *is* the option with the maximal expected choiceworthiness when the agent assigns little or no weight to her prudential reasons. And it permits an agent to donate her kidney to her dying sibling, even though this goes against her prudential reasons. This is because donating her kidney is the option with the maximal expected choiceworthiness when the agent assigns little or no weight to her prudential reasons.

This proposed way of reformulating MEC is not ad hoc. As I explained in Section Four, we generally think that there is a moral asymmetry regarding what an agent is permitted to do to herself and what she is permitted to do to others. Expanded MEC fails to respect this moral

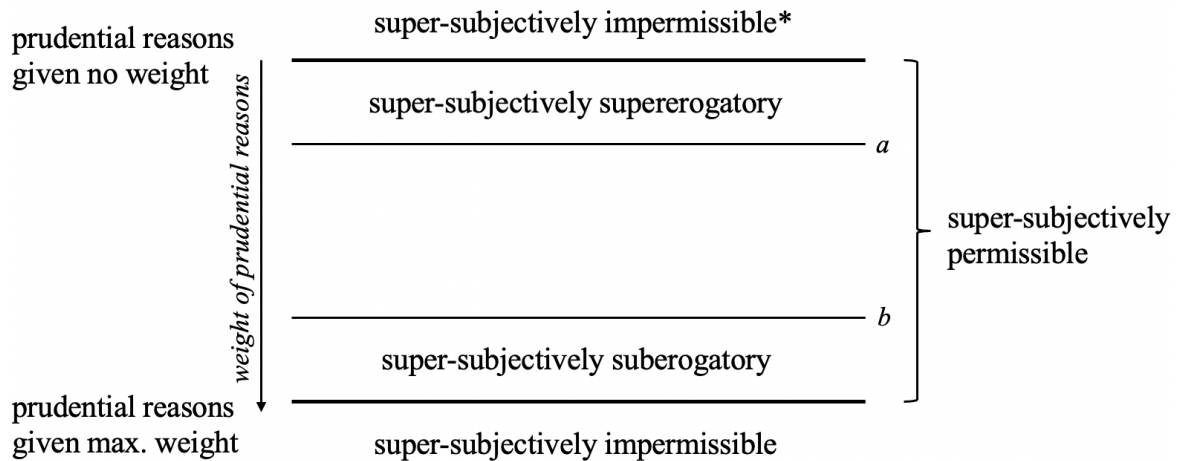
asymmetry, and that is why it faces the problem of suboptimal supererogation and the problem of optimal suberogation. Discretionary MEC, however, does respect this moral asymmetry, and that is why it avoids these problems. Furthermore, there are theoretical grounds to establish this kind of first-personal authority over the moral weight of one's interests. For instance, we can appeal to the principle that Jonathan Parry calls the "Power of Prudential Exclusion", which he uses to argue that we have the power to prevent someone from justifying their actions by appealing to the fact that they will benefit us, by repudiating those benefits (2017, 371). Presumably, this power to repudiate our own interests would likewise apply when it comes to determining what we ourselves ought to do. Similarly, Seth Lazar also argues that 'agent-centred options to favour and sacrifice one's own interests is grounded in a particular aspect of self-ownership' (2019, 36). These ideas show that Discretionary MEC is an independently plausible response to the problems with demandingness, supererogation, and suberogation. It respects the agent's authority over the moral weight of her own interests, permitting but not requiring the agent to exclude her interests entirely when determining what to do under moral uncertainty.

7. Going Beyond the Demands of Morality

So far, we've seen that the original version of MEC faces a dilemma: If it is sensitive only to moral reasons, and excludes prudential reasons, it is too demanding, as it requires us to follow theories which involve great personal sacrifice over common sense morality. However, if MEC includes prudential reasons, by turning to the all-things-considered choiceworthiness ordering of options, it results in counter-intuitive implications when it comes to our common-sense intuitions about supererogatory and suberogatory acts. We've also seen that MEC can be revised in a way that allows it to avoid this dilemma. These revisions give us Discretionary MEC.

There are, however, still some actions which are not handled correctly by Discretionary MEC because they remain super-subjectively impermissible. We said that according to Discretionary MEC, an act is super-subjectively permissible if and only if it has the maximal expected choiceworthiness when the agent gives her prudential reasons full, partial, or no weight. So, if an act does *not* have the maximal expected choiceworthiness even when the agent gives full weight to her prudential reasons, the act would be super-subjectively impermissible because it is suboptimal under Discretionary MEC. There is, however, another way that an act

may be suboptimal under Discretionary MEC. If an act does not have the maximal expected choiceworthiness even when the agent gives *no weight* to her prudential reasons, the act would also be super-subjectively impermissible. The diagram below demonstrates this:



Discretionary MEC, then, results in counter-intuitive implications when it comes to our common-sense intuitions about a certain category of supererogatory acts. Acts of self-sacrifice will remain suboptimal under Discretionary MEC if the moral agent has some amount of credence in moral theories which deem these acts to be impermissible. These intuitively supererogatory acts will not have the maximal expected choiceworthiness even when the moral agent gives no weight to her prudential reasons. Consider, for example, the following case:

*Burning Building**. A person is trapped inside a burning building. Fire-fighters are on their way, but you know that by the time they arrive, although they will be able to save the person's life, he will suffer from serious injuries, including the loss of both legs. You can enter the building to save the person from this fate, but it is certain that, as a result of doing so, you will suffer injuries even more serious than the loss of both legs.

Once again, suppose that there are two possible options: SELF-SACRIFICE and DO NOTHING. And once again, suppose some decision maker has high credence that both of these options are permissible, with SELF-SACRIFICE being supererogatory but not required. This time, however, suppose she also has some small credence in a moral theory like act utilitarianism according to which DO NOTHING is the only permissible option, because it is the option that produces least total harm. Even if the moral agent excludes her prudential reasons entirely, SELF-SACRIFICE will still be suboptimal in terms of choiceworthiness,

making it super-subjectively impermissible under Discretionary MEC. The option with the maximal expected choiceworthiness, even when the agent gives no weight to her prudential reasons, would be DO NOTHING, making it super-subjectively obligatory according to Discretionary MEC. So, even if we are almost certain that SELF-SACRIFICE is supererogatory and therefore morally permissible, so long as we have some amount of credence in first-order moral theories like act utilitarianism according to which these acts are impermissible, we would be super-subjectively prohibited from choosing these options even under Discretionary MEC.¹³

One might argue that this is a problem for theories like act utilitarianism, rather than for Discretionary MEC, because these theories fail to accommodate the common-sense thought that we are permitted to perform acts which result in less good overall, so long as the negative effects are experienced solely by the agent herself. So, proponents of Discretionary MEC (and MEC in general) might just say that this is the correct implication, given the agent's credences—that agents are super-subjectively prohibited from self-sacrificial deeds which do not promote the overall good, as these acts would be neither required by any moral theory, nor permitted by consequentialist theories.

However, just as the demandingness objection is a problem for MEC because we believe that what we super-subjectively ought to do is less demanding than what MEC implies, failing to accommodate for such supererogatory acts is problematic for Discretionary MEC because we believe that what we are super-subjectively permitted to do ought to be less restricting. With Discretionary MEC, it seems we are prohibited from performing acts which we are almost certain are supererogatory, and when doing so will involve no costs except to the moral agent herself. We generally regard such sacrificial acts as heroic, or saintly, or praiseworthy, and so such acts should at least be permissible under moral uncertainty.

I think there are a few ways to respond to this line of objection.

First, we can modify Discretionary MEC to capture our intuitions about such cases of supererogation by allowing the agent to give *negative* weight to her prudential reasons. If the agent gives negative weight to her prudential reasons when determining what to do under moral

¹³ This is all assuming that we take both options to be equally choiceworthy because they are both morally permissible. If, instead, we regard the supererogatory act as being *more* choiceworthy than the merely permissible act, it would not be a forgone conclusion that the option of SELF-SACRIFICE is super-subjectively permissible. However, if we take this position, it results in other problems for MEC, as I explain in Section 3. I thank an anonymous reviewer for pressing me on this.

uncertainty, then she would be permitted to perform supererogatory acts which result in a decrease in overall goodness, so long as the decrease in welfare is taken on by the agent herself. This solution, however, might well seem ad hoc. Also, it is doubtful whether the negative weight assigned to the agent's prudential reasons could offset the impermissibility of the act considering her credence in act utilitarianism, however small.

Another way to respond to the objection is to argue that it is plausible that we should have some credence in a moral theory on which we are morally required to perform acts of self-sacrifice for the sake of a lesser overall outcome. For instance, we could have some credence in a view like *Selfless Utilitarianism* which does not count the pleasure and pains that accrue to the moral agent herself, but only the pleasures and pains that accrue to others.¹⁴ If we have some amount of credence in a moral theory like this, this might offset the impermissibility of the self-sacrificial act under act utilitarianism, allowing such acts to be super-subjectively permissible under Discretionary MEC.

While moral theories like *Selfless Utilitarianism* seem somewhat implausible, it also seems to me that, if the moral agent gives act utilitarianism the benefit of the doubt despite being almost sure that acts of self-sacrifice are supererogatory, then they also should have some non-zero credence in views like *Selfless Utilitarianism*. Also, because the agent's credence in a theory like act utilitarianism is low, the agent need have only a small amount of credence in moral theories like *Selfless Utilitarianism* to counterbalance the impermissibility of the act under act utilitarianism. And depending on the specifics of the situation, the amount of credence in a theory like *Selfless Utilitarianism* needed to do the offsetting may be much smaller than her credence in act utilitarianism. For example, say that you can sacrifice 6 units of your well-being to provide 5 unit of well-being to a stranger. If you choose to SELF-SACRIFICE rather than DO NOTHING, this will result in a net loss of -1, but +5 if we count only the pleasures and pains that accrue to others. So, it seems that acts of self-sacrifice of this kind would be super-subjectively permissible even when you have less credence in *Selfless Utilitarianism* than in act utilitarianism. If, instead, you can sacrifice 6 units of your well-being to provide just 1 unit of well-being to a stranger, choosing SELF-SACRIFICE will result in a net loss of -5

¹⁴ See Sider (1993). Sider doesn't endorse *Selfless Utilitarianism* but a view he calls *Self/Other Utilitarianism*, which seems much more plausible. Having credence in a view like *Self/Other Utilitarianism*, however, will not be sufficient to solve the problem raised for Discretionary MEC, because it makes self-sacrificial acts which result in a lesser overall welfare merely morally permissible but not more choice-worthy.

under act utilitarianism and only +1 under *Selfless Utilitarianism*. Acts like this *would* be prohibited under Discretionary MEC, but that doesn't seem to be so implausible.

Finally, we could just accept that Discretionary MEC has its limitations and emphasise that it is at least less demanding or restricting than other views on the table. Discretionary MEC is considerably less demanding than the original version of MEC, and the range of supererogatory acts that Discretionary MEC prohibits is significantly smaller than those that Expanded MEC prohibits. Discretionary MEC seems to be the view that offers the best balance of theoretical costs and benefits.¹⁵

8. Conclusion

I have presented a new challenge to the all-things-considered version of MEC. I argued that, although appealing to the all-things-considered choiceworthiness ordering of options allows MEC to escape the demandingness problem, it gives rise to another problem: acts which we consider to be supererogatory are rendered suboptimal and hence super-subjectively impermissible, and acts which we consider to be suberogatory are rendered optimal and hence super-subjectively obligatory. This problem arises because when we factor in prudential reasons, we can be obligated to act in accordance with our interests and prohibited from acting against our interests.

We can solve these problems of suboptimal supererogation and optimal suberogation by revising MEC so that agents can choose how much weight to give to their prudential reasons, up to a limit. As we generally think agents are permitted to act against their interests, it should be up to the moral agent to choose how much weight to assign to their own interests, if any at all. So, we can say that, so long as an act is maximally choiceworthy either fully including, partially including, or excluding prudential reasons, it is super-subjectively permissible. This revision, which I call Discretionary MEC, allows us to maintain both the permissibility and optionality of supererogation and suberogation, while also capturing the way in which supererogatory acts are morally better and suberogatory acts are morally worse than other permissible acts.

In the final section, I showed that even with this new formulation, MEC might not be able to accommodate our intuitions regarding the permissibility of supererogatory acts which

¹⁵ See Kaczmarek and Lloyd (forthcoming) for a rival view to Discretionary MEC.

result in a net loss in overall good. I then suggested some ways in which we could overcome this limitation. But even with this limitation, Discretionary MEC seems to have the best balance of theoretical costs and benefits when compared with the original version of MEC and the all-things-considered version of MEC, and so it should be preferred to these other views.¹⁶

References

- Barry, Christian, and Patrick Tomlin. 2016. "Moral uncertainty and permissibility: evaluating option sets." *Canadian Journal of Philosophy* 46 (6): 898–923. <https://doi.org/10.1080/00455091.2016.1198198>
- Driver, Julia. 1993. "The Suberogatory." *Australasian Journal of Philosophy* 70 (3): 286–295. <https://doi.org/10.1080/00048409212345181>
- Gracely, Edward J. 1996. "On the Noncomparability of Judgements Made by Different Ethical Theories." *Metaphilosophy* 27 (3): 327–332. <https://doi.org/10.1111/j.1467-9973.1996.tb00212.x>
- Gustafsson, Johan E. and Olle Torpman. 2014. "In Defence of My Favorite Theory." *Pacific Philosophical Quarterly* 95 (2): 159–174. <https://doi.org/10.1111/papq.12022>
- Harman, Elizabeth. 2014. "The Irrelevance of Moral Uncertainty." *Oxford Studies in Metaethics* 10: 53–79. <https://doi.org/10.1093/acprof:oso/9780198738695.003.0003>
- Hedden, Brian. 2016. "Does MITE Make Right?" *Oxford Studies in Metaethics* 11: 102–128. <https://doi.org/10.1093/acprof:oso/9780198784647.003.0005>
- Kaczmarek, Patrick & Lloyd, Harry R. Forthcoming. "Moral Uncertainty, Pure Justifiers, and Agent-Centred Options." *Australasian Journal of Philosophy*.
- Lazar, Seth. 2019. "Self-Ownership and Agent-Centred Options". *Social Philosophy and Policy* 36 (2): 36–50. <https://doi:10.1017/S026505219000463>
- Lockhart, Ted. 2000. *Moral Uncertainty and its Consequences*. New York: Oxford University Press. <https://doi.org/10.1093/mind/111.443.693>

¹⁶ I am especially grateful to Joe Horton for invaluable feedback on multiple drafts, and Nikhil Venkatesh for helpful written comments on a very early version of the paper. I also thank two anonymous reviewers for the *Canadian Journal of Philosophy*, whose comments vastly improved the paper. For engaging discussions and constructive feedback, I am indebted to the group of brilliant philosophers at the Global Priorities Institute (Oxford) and the audiences of the Uncertain Oughts workshop (Helsinki), the UCL-Leipzig Graduate Conference, and the Open Minds Graduate Conference (Manchester).

- MacAskill, William, Krister Bykvist & Toby Ord. 2020. *Moral Uncertainty*. Oxford: Oxford University Press. <https://doi.org/10.1093/oso/9780198722274.001.0001>
- MacAskill, William and Ord, Toby. 2018. ‘Why Maximize Expected Choiceworthiness?’, *Noûs* 54 (2): 327–353. <https://doi.org/10.1111/nous.12264>
- Muñoz, Daniel & Baron-Schmitt, Nathaniel. 2024. ‘Wronging Oneself’, *Journal of Philosophy* 121 (4): 181-207. <https://doi.org/10.5840/jphil2024121416>
- Parfit, Derek. 2011. Moral Concepts, in *On What Matters*, vol. 1, edited by Samuel Scheffler, Oxford: Oxford University Press, 150–174. <https://doi.org/10.1093/acprof:osobl/9780199572809.001.0001>
- Parry, Jonathan. 2017. “Defensive Harm, Consent, and Intervention.” *Philosophy & Public Affairs* 45 (4): 356–396. <https://doi.org/10.1111/papa.12099>
- Rosenthal, Chelsea. 2020. “What decision theory can’t tell us about moral uncertainty”. *Philosophical Studies* 178 (10): 3085-3105. <https://doi.org/10.1007/s11098-020-01571-3>
- Ross, Jacob. 2006. “Rejecting Ethical Deflationism.” *Ethics* 116 (4): 742–768. <https://doi.org/10.1086/505234>
- Sepielli, Andrew. 2009. “What to do when you don’t know what to do.” *Oxford Studies in Metaethics* 4: 5–28. <https://doi.org/10.1093/oso/9780199566303.003.0002>
- Sepielli, Andrew. 2018. Decision-Making Under Moral Uncertainty, in Karen Jones, Mark C. Timmons & Aaron Zimmerman (eds.), *Routledge Handbook on Moral Epistemology*. New York: Routledge. <https://doi.org/10.4324/9781315719696>
- Sider, Ted. 1993. “Asymmetry and Self-Sacrifice.” *Philosophical Studies* 70 (2): 117–132. <https://doi.org/10.1007/BF00989586>
- Slote, Michael. 1984. “Morality and self-other asymmetry.” *Journal of Philosophy* 81 (4): 179-192. <https://doi.org/10.2307/2026119>
- Stocker, Michael. 1976. “Agent and Other: Against Ethical Universalism.” *Australasian Journal of Philosophy* 54 (3): 206–220. <https://doi.org/10.1080/00048407612341221>
- Urmson, James O. 1958. Saints and Heroes, in *Essays in Moral Philosophy*, ed. A. I. Meldon. University of Washington Press.
- Weatherson, Brian. 2019. Normative Externalism. Oxford: Oxford University Press.
- Weatherson, Brian. 2002. “Review of Ted Lockhart’s ‘Moral Uncertainty and its Consequences.’” *Mind* 111 (443): 693–96. <https://doi.org/10.1093/mind/111.443.693>