

This paper is forthcoming in *American Philosophical Quarterly*. This is a draft copy – please cite the published version where possible.

Supererogation and the Case Against an ‘Overall Ought’

Introduction

Consider Helen.

Helen accepts that, given the needs of the poor, she is morally justified in keeping for herself only what she requires to meet her basic needs. She also thinks it’s important to do what is right, but she really wants to go hiking amidst spectacular mountain scenery, which involves spending money on travel, accommodation, and hiking equipment. She doesn’t think that her desire to go hiking provides a moral justification for spending this money, but she also doesn’t think that it is irrational for her to spend it.¹

Helen concludes that what she morally ought to do is to give most of her money to charity, and what she ought to do based on her own self-interest is to spend it all on hiking. Helen might want to ask a further question: what ought she do overall, taking these different reasons into account?

This paper has two aims. Firstly, it will argue that the ‘overall ought’ is not treated carefully enough in philosophical discourse, and is in need of clarification. The second aim is more ambitious: to give a new argument for why the targeted overall ought does not stand up to scrutiny.

¹ (Singer, 2009 p.389)

The main problem I'll pose for my opponent will be a version of the paradox of supererogation,² and I'll also explain why it's only a problem for this specific kind of overall ought, rather than for non-moral oughts generally (such as prudential oughts). The problem, briefly stated, is this: when our overall obligations and our moral obligations come apart, then the overall ought obligates us *not* to follow our moral obligations. This, I will argue, is implausible. I will demonstrate that for four different analyses of what the 'overall ought' might mean, they will either come across a form of this problem or will not closely enough resemble the kind of concept people refer to when they talk about what we overall ought to do. First, I will explain the target concept in more detail.

1. What We Overall Ought to Do

I take the kind of overall ought the paper targets to be widely used. Dancy, for example, discusses an 'overall ought' as what we have most reason to do given all the 'contributory reasons',³ and Davidson talks about what we should do all-things-considered.⁴ Another example is Hurka, who describes the overall ought as it is used by a school of thought he calls the 'Sidgwick-Ewing' school.⁵ Zimmerman discusses an overall ought from the perspective of virtue ethics.⁶ But one of the best treatments of the overall ought can be found in McLeod,⁷ as he argues for a coherent theory of an overall ought (which he refers to as the 'Just Plain Ought'). He frames it in terms of the kind of question above, which asks what to do when there are conflicting oughts. He then says,

² The paradox of supererogation has been discussed before, but not specifically why it's a problem for 'overall ought' theorists. The problem generally is the tension between two claims about supererogatory acts. Firstly, there's the claim that supererogatory acts are more than what's required of agents. Secondly, there's the claim that we ought to do what's morally best, and therefore we are required to do what's supererogatory. For example, Heyd refers to this phenomenon as the 'good-ought tie-up' in (Heyd, 2016). See also (Dancy, 1988), (Williams, 2011 p.198-199) and (Portmore, 2003).

³ (Dancy, 2003)

⁴ (Davidson, 1970)

⁵ (Hurka, 2014)

⁶ (Zimmerman, 2008)

⁷ (McLeod, 2001)

The [overall ought] is the idea of an “ought” that is not identical to any of the relative or qualified “oughts” – that is, the moral “ought,” the prudential “ought,” the aesthetic “ought,” and so on. [...] ...the concept of [the overall ought] is distinct from any relative “ought” concept.⁸

I’ll try to clarify this idea.

This paper’s target is a particular concept, one that people refer to when they talk about what they *overall ought* to do. There are multiple ways that people might use the term, but the concept this paper will focus on is a concept of what we ‘overall ought’ to do that (i) tries to find a balance between different kinds of reasons⁹ (such as moral and prudential) and (ii) does so by appealing to an overall standard. My opponents do not always explicitly mention these desiderata, but I have laid them out to help define exactly which concept I am targeting.

Turning first to (i), Helen, introduced above, has moral reasons and prudential reasons, which we can assume come apart.¹⁰ Similarly, the target theory assumes that moral reasons are not always overriding,¹¹ that sometimes an agent is not required¹¹ to make personal sacrifices in order to do what’s morally best. It appears when we think we *should* find a balance between, say, helping others and looking after our own interests.¹² According to my opponent, there’s a need for

⁸ (McLeod, 2001 p.273)

⁹ I use ‘reasons’ in this paper to mean considerations that count in favour of an act or group of acts. See the ‘primitive’ description given by, for example, (Scanlon, 1998 p.17).

¹⁰ You might think that prudential and moral oughts never come apart. Anscombe and later-Foot, for example, can be seen as holding that view in (Anscombe, 1981) and (Foot, 2001), and those who want to look even further back can find it in the likes of Aristotle in, for example, (Aristotle, 2009). These accounts are not the target of the criticisms in this paper.

¹¹ My opponent might think that moral *obligations* or *oughts* are overriding, but not that moral reasons or the morally best actions will always result in these. Stroud seems to use discussion of an ‘overall ought’ in her paper on moral overridingness, (Stroud, 1998)

¹² You might want to find such a balance because of concerns about moral demandingness, for example. That is, if you worry that doing the maximally best thing in a situation is too demanding, then you might appeal to a different kind of ‘ought’ that lets you find a balance between that demanding morally best option and options that take *your own* interests into consideration as well. See, for example, (Scheffler, 1992), (Benn, 2015) and (McElwee 2016) on demandingness objections, (Scheffler, 1994) and (Norcross, 2006) for the demandingness of consequentialism, (Baron, 1987) and (Annas, 1984) for discussion on demandingness in Kantian ethics and (Ashford, 2003) for discussion of demandingness in Scanlon’s contractualism.

something *more* to the story than the reasons themselves, something that tells us whether Helen is balancing her other reasons *correctly*, whether she's giving the *right* weight to moral versus non-moral reasons. Woollard, for example, expresses such a need in her review of Singer's original example of Helen.¹³ She wants to ask a further question: what ought Helen do overall? What's the correct amount of consideration to give to these different kinds of reason?

Moral reasons and prudential reasons are two examples of what you might want to balance, but the list might also include things like hiking reasons, Catholic reasons or bearded-dragon-owner reasons.¹⁴ That is, reasons you might have to promote or respect the values of hiking, Catholicism, or looking after your bearded dragon(s).¹⁵ The 'overall ought' might be just a balance of the prudential and the moral, or it might, instead, be a balance of many other kinds of reason. And there will be overlap within these categories, since to a large extent they're artificial. Because Helen enjoys hiking, for example, then what we'd call her hiking reasons and prudential reasons will often overlap.

A good analogy here might be with assessing a film. When asked to choose a favourite film, or to rank a selection, someone might have some specific criterion in mind that they judge it on. I might, for example, pick the film *March of the Penguins*, based on the criterion of 'the film which has the most penguins'. But, perhaps more regularly, I might try to find a film which is not just the best at representing a high number of penguins but the best *overall*, given multiple criteria that I care about: lighting, good direction, theme, *and* number of penguins. (As it happens, this would still lead me to pick *March of the Penguins*.)

¹³ (Woollard, 2010)

¹⁴ A similar list is discussed in, for example, (Broome, 2007) where he discusses normative 'requirements', which I take to be similar to the kinds of 'oughts' I'm discussing.

¹⁵ The language of promoting and respecting is taken from (Pettit, 1993). I use it here to show the account's neutrality in regard to different approaches to value.

Next, (ii) was the idea that the overall ought is appealing to an ‘overall’ standard. I don’t have much to say about what this overall standard is, since I find this unclear myself. It may well be something we just have to work out through reasoning, using our judgment.¹⁶

For the rest of this paper I’ll use abbreviations for three actions that Helen could possibly take when deciding how to spend her time and money:

(M) The action that Helen *morally* ought to do.

(P) The action that Helen *prudentially* ought to do.

(O) The action that Helen *overall* ought to do.

We can suppose that the action denoted by (M) won’t (unless circumstances are particularly unusual) involve any hiking for Helen, rather it will most likely involve her giving away a large proportion of time and effort to charitable causes. For the sake of making this less complicated, let’s also suppose what would actually be the best for Helen prudentially doesn’t *much* overlap with what would be morally best for her to do, rather it would involve saving up for and having a peaceful but resource-consuming hiking career.¹⁷ Then (O) would be the action prescribed by the ‘overall ought’, the one which is the ‘right’ balance between the other two.

Now I’ve tried to clarify the concept I’ll list four possible ways to analyse what it may mean in moral discourse:

(1) The overall ought tells an agent what is demanded of them, what they are obligated to do.

(2) The overall ought tells an agent what it is praiseworthy for them to do.

(3) The overall ought tells an agent what they have most reason to do.

¹⁶ This is the kind of reasoning that Crisp refers to in (Crisp, 1996). It might also be what McLeod thinks the overall ought means in (McLeod, 2001) when he talks about it being a standpoint of ‘reason’ or ‘reason-as-such’, and what McElwee argues we should interpret it as in (McElwee, 2007 p.366).

¹⁷ To repeat footnote 11: those who don’t see the moral and the overall oughts as ever coming apart are not the targets of this paper’s argument.

(4) The overall ought tells an agent what the minimum that they are obligated to do is.

I've listed four possibilities here,¹⁸ and I'll tackle them one at a time. Some of these may often overlap, such as (1) and (3), and I'll discuss that in more detail in their individual sections later. Some of these analyses may also sound more natural than others. To me, the most natural interpretation of the overall ought is (1), but after I've argued against it my opponent may want to retreat to one or more of the other options. Options (2)-(4) may seem more plausible when (1) has been ruled out. Because of this, it is important to show why every one of these analyses is problematic. In the end I will argue that the analyses which are the least vulnerable to my criticisms on the grounds of the problem of supererogation will be the *most* vulnerable to the charge that they do not sound like an 'overall ought' after all.

2. The Problem of Supererogation and the Overall Ought

For each of the four analyses I will argue that the overall ought cannot work. In some cases this will be because accepting that analysis of the overall ought will also mean accepting some implausible claims about supererogatory acts. In other cases, the overall ought will have had to be qualified so much (to avoid the first problem) that it no longer resembles the kind of 'overall ought' that this paper targets. At a very minimum, by the end of this paper the spotlight will be on my opponent to clarify their position and explain what route they're going to use to navigate away from these problems.

¹⁸ Bart Streumer in (Streumer (2007) p.354) gives a different list of four interpretations of ought claims generally, although not specifically of the overall ought. As well as an ought meaning that the agent might have an obligation (which I've covered with (1)) and that it would be what the agent has 'most reason' to do (which I've covered with both (1) and more specifically with (3)) he also lists the interpretation that the agent might just be expected to act in that way, and the interpretation that it would be good if the agent acted in that way. I have no problem with either of those last two interpretations of the overall ought, but take them to be different from the kind that this paper is arguing against.

First, I will briefly introduce the concept of supererogation in a little more detail.¹⁹ The concept can be traced back to Urmson, who argued that moral theories needed to be able to account for a category of actions that were ‘saintly’ or ‘heroic’.²⁰ Such supererogatory actions are those which are morally good (and usually exceptionally so) but are either not required or at least not *morally* required.²¹ Those who agree that some acts are supererogatory might think that in Helen’s case, action (M) is supererogatory, and that it would be good but not required of her to dedicate all of her resources to others.

This paper is not committed to the existence of supererogatory acts. Rather, supererogatory acts are something that my opponent is committed to. This is because the concept of the overall ought I’m working with relies on the overall ought sometimes differing from the moral ought, and that would mean that the morally best option is not required of us in at least one sense.²²

I’ll now turn to the four different analyses of what someone might mean when they use an overall ought.

(1) Demand and obligation

The overall ought tells an agent what is demanded of them, what they are obligated to do.

One job for the overall ought may be to describe what it is that is demanded of an agent, what the agent is obligated to do. This may be the most immediately natural-sounding interpretation of

¹⁹ For some helpful discussion on supererogation generally, see, for example, (Archer, 2016), (Benn, forthcoming), (Horgan and Timmons, 2010) and (McElwee, 2017).

²⁰ (Urmson, 1958 p.199)

²¹ Archer refers to these conditions as ‘Morally Optional’ and ‘Morally Better’ in (Archer, 2016). He talks about the acts not being *morally* required, whereas other sources like (Heyd, 2016) talk about the acts not being “(strictly) required”. For the purposes of this paper I take both kinds to be cases of supererogation.

²² One might still hold that the morally best option is *morally* required, even if those moral requirements are not overriding, and they might take that to be the case even if what’s *morally required* of us is not what we overall ought to do. Even given this distinction, the morally best act is still supererogatory in the sense that it is ‘above and beyond’ what we overall ought to do, and that’s enough for the purposes of this paper.

overall ought language. After all, it tells the agent what they *ought to do* overall, so it's understandable to think that this might consist in some kind of normative obligation or demand.²³

Under this interpretation, (O) represents what Helen is obligated to do all-things-considered; it takes into consideration not just her moral reasons to donate money to charity but it also considers how much weight she should give to these reasons versus her prudential reasons and determines the correct middle-ground between them.

The problem with looking at the overall ought as what an agent is *obligated* to do is that it means that Helen is overall *obligated* not to give up *more* of her resources to charity than whatever amount (O) would involve. Giving up more than (O) requires of her is not just something she is *not obligated* to do but something that she is *obligated not to do*. Helen's moral reasons have been weighed against her desire to go hiking, and the balance of what she is required to do has been found. But it sounds very implausible to say that Helen is obligated not to do the supererogatory act. Intuitively, it would be better if Helen gave away more of her wealth. Indeed, we call it an act of supererogation precisely because it is such a great thing for someone to do.

My opponent might claim that this objection confuses what's morally better and what's overall better. Sure, they might reply, it's not plausible to describe a morally supererogatory agent as being *morally* worse, but that's not what they would need to be committed to here. All they're saying is that the behaviour of the agent in question is 'overall' worse. We might think it sounds strange, but perhaps the strangeness is just because we're generally used to thinking of these terms as being moral terms. But the overall best is already defined as being different from the morally optimal in cases like Helen's. It would sound just as strange to say that the agent taking the best option is acting *prudentially* worse, but in this particular example that would still be true.

²³ There's a separate option according to which the overall ought is more like a minimum requirement of what an agent is obligated to do, but I'll address this separate concern under the umbrella of (4), since I think that this particular interpretation is more like (4) than it is like (1). For now, I'll take the overall ought to refer to the single act (insofar as it's possible to narrow it down that specifically) which we are obligated to do.

I don't think this is the kind of move that my opponent can plausibly make. This is because although they would want to deny that agents are required or obligated to perform supererogatory actions, I don't think that they would want to say that the agent is worse, even 'overall worse', for doing so. I would be surprised if my opponent is happy to condemn such a supererogatory agent for weighing her reasons incorrectly and placing *too much* emphasis on morality.

Part of the problem with the overall ought, and why it has so much trouble with supererogatory acts, is that it seems to act like a *quasi*-moral ought. It purports to tell us what's required, but while denying that it's always better to do what's morally best. Telling Helen that overall she should give x amount of her money to charity might, at first, seem like a good way to balance her reasons, but to describe this overall ought as something that tells her the *right* balance is not plausible.

This quasi-moral status of the overall ought is also why the overall ought, specifically, is the only kind of ought that comes across the paradox of supererogation in this way. Prudential oughts, bearded-dragon-owner oughts or hiking oughts, for example, aren't vulnerable to the paradox of supererogation. Prudential oughts tell an agent to do what's best for them, bearded-dragon-owner oughts tell an agent what best to do for their bearded dragon, and hiking oughts tell an agent how best to further the ends of hiking, and these are often not going to be the same as what's morally best. But these oughts don't disguise themselves as being anything other than what they are: oughts grounded in certain non-moral ends. To describe a morally supererogatory Helen as not fulfilling her prudential obligations doesn't seem implausible. To describe her as not fulfilling her overall obligations is uniquely problematic.

I have argued here that the overall ought is implausible under description (1), because it requires agents to not act supererogatorily. As I said earlier, (1) is possibly the most natural-sounding way to understand the overall ought. But now that I've raised an objection with it, my opponent may want to adopt a different understanding of the overall ought to try to avoid the problem. Some of these may seem less intuitively like 'oughts', but they may still be the most

charitable interpretation of what my opponents might mean with overall ought language, given the problems with (1). Next I'll look at three routes they might take. Firstly, they might want to retain an overall ought but without as much normative *oomph*, that is, without the same kind of demand / obligation on the agent. They could do this by understanding the overall ought as only being a description of what we have most reason to do, which I'll address in section (3) only after I've first set aside (2): the possibility that we analyse the overall ought in terms of praiseworthiness. I'll tackle the latter first because my response to it is so similar to my objection to (1), and so it follows more naturally. Then, in (4), I'll address a third potential escape route: understanding the overall ought not as an obligation to perform a single action, but as an obligation to perform one of a range of actions. This is analysis (4): as setting a minimum requirement for what agents are obligated to do.

(2) Praiseworthiness

The overall ought tells an agent what it is most praiseworthy for them to do.

We might take the overall ought to direct agents towards what it's *praiseworthy* for them to do. Suppose Helen is having trouble deciding what to do, and she turns to ask what she ought to do. "Overall, you ought to (O)" we might reply. It seems plausible to think that Helen isn't under an *obligation* to act in that way, but rather we're just telling her what option would be the most praiseworthy.

This comes across a similar problem to that raised in (1), and correspondingly my response will be shorter. I criticised the first analysis of the overall ought on the grounds that it entailed implausible claims about supererogatory acts: that agents were overall obligated *not* to perform them. The same criticism can be made of interpretation (2); it would describe (O) as being the most praiseworthy act, and (M) (or even some middle-ground between the two) would therefore be described as *less praiseworthy*. This is still implausible, and would make for a difficult bullet to bite.

Perhaps my opponent might find (2) appealing because of arguments like that of Wolf in ‘Moral Saints’. She argues that the kind of person who does perform supererogatory actions all the time is not really a very appealing kind of person either to be or to befriend.²⁴ This might give my opponent reason to think someone more balanced, likely to perform the overall actions like (O) rather than to go all out towards (M), might be more deserving of praise. But this isn’t so. Take this quote from Wolf,

Despite my claim that all-consuming moral saintliness is not a particularly healthy and desirable ideal, it seems perverse to insist that, were moral saints to exist, they would not, in their way, be remarkably noble and admirable figures. Despite my conviction that it is as rational and as good for a person to take Katherine Hepburn or Jane Austen as her role model as Mother Theresa, it would be absurd to deny that Mother Theresa is a morally better person.²⁵

The opponent who wants to use the overall ought to signal praiseworthiness would have to bite a bullet in which an idealised Mother Theresa is overall *less* praiseworthy for giving extra weight to her moral reasons. And although Wolf uses the phrase ‘morally better’ rather than ‘overall better’, we can see (and saw in (1)) that the point still stands.

Even if my opponent thinks that Wolf’s characterisation of moral saints is persuasive, this is *still* not enough to save this interpretation of the overall ought. My objection doesn’t rely on agents being complete moral saints like the ones Wolf describes. All I need for my objection to work is for an agent to be able to perform an even *slightly* morally better action than (O). Such an agent wouldn’t be subject to Wolf’s cutting verdict on the character of moral saints, but my

²⁴ (Wolf, 1982)

²⁵ (Wolf, 1982 p.432)

opponent would still have to accept that they are (slightly) less praiseworthy for not doing (O). This still seems implausible.

I've argued against interpretation (2) in a similar way to my argument against interpretation (1). Here my opponent has the same two options: accept implausible claims about agents who act supererogatorily, or not accept this interpretation of what the overall ought means.

(3) Most reason

The overall ought tells an agent what they have most reason to do.

Another suggestion for what the overall ought may indicate is that it might simply be what the agent has most reason to do. Although I have already defined the overall ought as a balance of an agent's reasons, analysis (3) understands it as only a *description* of what the agent has most reason to do, and is silent on whether we are also obligated to follow those reasons, as we saw with interpretation (1).

My criticism here begins with considering whether we can understand the overall ought as (3) without the obligation that came with (1), or if the two necessarily go together. Let's suppose first that (1) doesn't necessarily apply. Here my opponent comes across a new problem: (O) might be what an agent has most reason to do, but unless telling them so also comes with the kind of obligatory force discussed in (1) then it doesn't sound much like an overall ought. The kind of overall ought that I think my opponent is after, and the kind that I established at the beginning of this paper, needs to be something more than just a simple description of weighed reasons. That is, it needs to carry some kind of normative weight that to some extent *obligates* the agent to actually follow those reasons. Overall ought language is prescriptive, but without the normative force it's only *descriptive*. An understanding of the overall ought without being prescriptive tells the agent what the balance of their reasons is, but not that they should then do the thing that they have most

reason to do. And if my opponent chooses to interpret the overall ought as both (3) *and* (1), then it is subject to exactly the same problems already discussed.

There are two moves my opponent might want to make here. Firstly, they might argue that the normative force just comes from the fact that (O) is whatever the agent has most reason to do. People are just obligated to do what they have most reason to do, that's part of what it is to be a rational agent.²⁶ The agent has a rational reason to balance her prudential, moral and other reasons and so she does, and rationally she's then required to perform that action. But if my argument against (1) was convincing, then it still applies if we describe the overall ought as an ought that comes from rationality.²⁷ My opponent would have to accept that morally supererogatory agents are doing something rationally wrong, they're failing to fulfil their obligations to rationality, failing to do what they have most reason to do.

Suppose Helen gives away more of her resources to charity than (O) required of her. There would seem to be a tension in how we evaluate her action. We think that she has morally performed well, but rationally performed poorly. She has done what's right in terms of morality, and failed to do what's right according to rationality. Even though 'rightness' here tracks two different things, the way we evaluate rational actions and moral actions often coincides. If rationality and morality come apart, I'm not convinced that talk of what we 'overall ought to do' tracks the former. Of course, if my opponent does find this convincing, I'm happy at least to have clarified what the quasi-moral 'overall ought' really refers to.

The other move my opponent might want to make is to appeal to a normative force *less* than an obligation which directs the agent to follow their reasons. Because it is weaker than an obligation or a demand, my opponent might say, it escapes much of the force that comes with the

²⁶ For more detailed discussion on these kinds of questions of rationality, see for example Kieseewetter (2017)

²⁷ Wallace argues for a similar point: that "[w]e can perhaps say what it is rational to do with an eye to morality, and what it is rational to do with an eye to an individual's good, but there seems to be no common currency in terms of which to cash out claims about what it is most rational to do overall." (Wallace, (2006) p.131)

supererogation objection. Perhaps the overall ought *recommends* (O), rather than demands it, but in such a way that agents aren't condemnable for failing to do it. But this response fails, because the problem of supererogation comes with any level of normative force. It still seems implausible to overall criticise a moral saint on any level. Furthermore, this leads my opponent to a dilemma. The stronger the normative force to do what's overall best, the more obvious the problem of supererogation is. But the weaker the normative force is, the less like an 'overall ought' the overall ought sounds like. We use overall oughts to advise, to prescribe action, and that sounds most plausible with more normative force, more 'shouldness', behind it.

In this section I've argued that we don't use the 'overall ought' to mean what we have most reason to do without also including a kind of normative force that I argued against in (1). This time my opponent has two options. Firstly they could accept (3) in conjunction with (1) and, along with it, the implausible claims about supererogation. Their second option is to understand (3) on its own, as a simply descriptive claim that tells the agent facts about their reasons but not whether they should act on those reasons. This also seems implausible as a description of the overall ought.

(4) Minimum requirement

The overall ought tells an agent what the minimum that they are obligated to do is.

Finally I'll discuss the interpretation of the overall ought on which it doesn't pick out the only act that you're obligated to do, but picks out what the minimum act is that the agent ought to do. Here, we can understand it as Helen being advised that "overall, you ought to (O). You could donate more of your resources to charity than that, but (O) is the least you should do."




Figure 1

On this interpretation of the overall ought, the agent is obligated to do one of a range of permissible actions. They are permitted to do anything between (O) and (M). Let's examine how the range of permissible acts between (O) and (M) can, themselves, be interpreted. We have a range of acts that it is permissible (or 'overall allowed') for the agent to do, (see Figure 1) which range from (O) to (M). There are three possibilities:

- (a) The agent should do something more like (O), but everything in the range is permissible.
- (b) The agent should do something more like (M), but everything in the range is permissible.
- (c) All options on the scale have equal merit. The agent should do any option between (O) and (M), it doesn't (overall) matter which.

None of these three look promising. (a) is the easiest to dismiss since it stumbles across the same problem that we've already encountered in (1) and then seen repeated throughout multiple interpretations of the overall ought. As I've argued above, it's implausible for an 'overall' normative theory to prefer the agents who do not act supererogatorily over those who do. Option (a), in having a preference for acts like (O) rather than (M), comes across these same problems.

What about (b), in which it is overall preferable for the agent to perform acts that are more like (M)? This time we have a sliding scale that makes sense in terms of morally exemplary agents, since the more that the agent pays attention to her moral reasons the better the agent has done. The problem here is that it doesn't look much like the *overall* ought that my opponent started out with. Instead, it looks like a regular moral ought, with a line drawn at (O). Firstly, my opponent would need to accept that when they describe what someone like Helen overall ought to do, what they *actually* mean is that Helen ought to *at least* perform (O) but preferably pay even more attention

to her moral reasons. That is, when they say Helen overall ought to (O), they really mean it would be overall best if she did (M) instead. This doesn't sound like a plausible way to hear the description of what Helen overall ought to do.

Finally I'll turn to (c): the possibility that all options on the scale between (O) and (M) have overall equal merit. Again, this interpretation doesn't seem like the kind of thing people mean when they talk about what they overall ought to do, and this time it's because it doesn't actually direct the agent to act in a certain way. We might weigh up Helen's different options and advise her that she ought to allow herself the resources to go hiking once every couple of months. This doesn't sound like we're actually pointing her to a wide range of actions; it sounds like we've decided what the best option is, taking all of her reasons into account.

Interpretation (4) of the overall ought was the most complex, and gave my opponent several possible answers. It was the idea that the overall ought told the agent what was the minimum amount required of them, and any action between this and the morally best option would be permissible. With this in mind, I separated three ways to understand that claim into (a), (b) and (c). If my opponent wanted to accept (a), they would have to accept the implausible claims about supererogation that have been the basis of my main argument against the overall ought. For (b), my opponent would also need to accept that the overall ought is just a moral ought, and that when they tell an agent what she must overall do, they actually mean it would be overall better for them to do something *other than* the act they've just prescribed; better for the agent to do (M) than (O). Finally, for (c), my opponent would have to accept that when we tell an agent what they overall ought to do then we're not actually telling them to do that particular act, but rather a larger range of options. Not very informative after all. It seems, then, that no analysis of this kind of 'overall ought' language is unproblematic.

Conclusion

The target of this paper was a quasi-moral overall ought that aims to tell us how to balance our prudential and moral reasons. By detailing the range of things that people can mean by an overall ought I've made the case that (at the very least) such language needs to be used more carefully and explained more thoroughly in philosophical discourse.

This paper also argued that commitment to such an overall ought is not plausible. The overall ought theorist, I've argued, runs into trouble when it comes to describing supererogatory agents. One route my opponent might have taken was to understand the overall ought as carrying less normative force, and so avoiding the objection this way. But the less force the overall ought carries, the more other problems stand out, such as not being able to do the job they it's supposed to, to tell Helen how much consideration to pay to her moral reasons compared to her prudential reasons.

To what extent should moral philosophers be worried? I'll give a few final remarks here on the implications of losing a certain kind of overall ought. McLeod argued that the overall ought was important for a number of discussions in moral philosophy, including those about our ability to reason practically and whether we should obey moral oughts.²⁸ This paper's argument will have an effect on those questions, but not leave them unanswerable. There may be other ways to use overall ought language that isn't covered by the criticisms in this essay, and that aren't an appeal to the kind of problematic, quasi-moral, 'overall' standard that I've argued against. Perhaps, for example, by using more concrete measurements of correctness, such as through appealing to the agent's desires or certain social standards. An exploration of such issues will have to wait until another time.

BIBLIOGRAPHY

²⁸ (McLeod, 2001 p.274-249)

- Annas, J. (1984) 'Personal Love and Kantian Ethics in *Effi Brieset*' in *Philosophy and Literature* pp.15-31
- Anscombe, G. (1981) 'Modern Moral Philosophy' in *Ethics, Religion and Politics: Collected Philosophical Papers Volume III*, Oxford: Basil Blackwell
- Archer, A. (2016) 'Supererogation, Sacrifice and the Limits of Duty' in *Southern Journal of Philosophy* vol.54 no.3
- Aristotle (2009) *The Nichomachean Ethics: Oxford World's Classics*, Oxford: Oxford University Press
- Ashford, E. (2003) 'The Demandingness of Scanlon's Contractualism' in *Ethics*, no.2 pp.273-302
- Baron, M. (1987) 'Kantian Ethics and Supererogation' in *Journal of Philosophy*, vol.84 no.5 pp.237-262
- Benn, C. (forthcoming) 'Supererogation, Optionality and Cost' in *Philosophical Studies*
- Benn, C. (2015) 'Over-Demandingness Objections and Supererogation' in *The Limits of Moral Obligation*, Marcel van Ackeren and Michael Kühler (ed.) New York: Routledge
- Broome, J. (2007) *Requirements*, available online at:
<http://www.fil.lu.se/hommageawlodek/site/papper/BroomeJohn.pdf>
- Crisp, R. (1996) 'The Duality of Practical Reason' in *Proceedings of the Aristotelian Society* vol.96 no.1 pp53-73
- Dancy, J. (2003) 'What Do Reasons Do?' in *Southern Journal of Philosophy* vol.XLI
- Dancy, J. (1988) 'Supererogation and Moral Realism' in *Human Agency*, J. Dancy and others (ed.) Stanford: Stanford University Press
- Davidson, D. (1970) 'How is Weakness of the Will Possible?' in *Moral Concepts*, Joel Feinburg (ed.) Oxford: Oxford University Press
- Foot, P. (2001) *Natural Goodness*, Oxford: Oxford University Press
- Heyd, D. (2016) 'Supererogation' in *Stanford Encyclopedia of Philosophy Spring 2016 Edition*, Edward N. Zalta (ed.) Available online at:
<http://plato.stanford.edu/archives/spr2016/entries/supererogation/>
- Horgan, T. and Timmons, M. (2010) 'Untying a Knot from the Inside Out: Reflections on the Paradox of Supererogation', in *Social Philosophy and Policy*, vol.27 issue 2, pp.29-63
- Hurka, T. (2014) *British Ethical Theorists from Sidgwick to Ewing*, Oxford: Oxford University Press

- Kiesewetter, B. (2017) *The Normativity of Rationality*, Oxford: Oxford University Press
- McLeod, O. (2001) 'Just Plain Ought' in *Journal of Ethics*, vol.5 no.4, pp.269-291
- McElwee, B. (2017) 'Supererogation Across Normative Domains' in *Australasian Journal of Philosophy*, vol.95 issue 3, pp.505-516
- McElwee, B. (2016) 'Demandingness Objections in Ethics' in *The Philosophical Quarterly* 67 pp.84-105
- McElwee, B. (2007) 'Consequentialism, Demandingness and the Monism of Practical Reason' in *Proceedings of the Aristotelian Society* 107 pp.359-374
- Norcross, A. (2006) 'The Scalar Approach to Utilitarianism' in West, H. (ed.) *The Blackwell Guide to Mill's Utilitarianism*, Oxford: Blackwell Publishing
- Pettit, P. (1993) 'Consequentialism' in P. Singer (ed.) *A Companion to Ethics*, Oxford: Blackwell Publishing Ltd.
- Portmore, D. (2003) 'Position-Relative Consequentialism, Agent-Centred Options and Supererogation', *Ethics* 113 pp.303-322
- Scanlon, T. (1998) *What We Owe To Each Other*, London: Harvard University Press
- Scheffler, S. (1992) *Human Morality* Oxford: Oxford University Press
- Scheffler, S. (1994) *The Rejection of Consequentialism*, Oxford: Clarendon Press
- Singer, P. (2009) 'Reply to Michael Huemer', in J. Schaler (ed.) *Peter Singer Under Fire*, Chicago and La Salle: Open Court Publishing Company
- Streumer, B. (2007) 'Reasons and Impossibility' in *Philosophical Studies* vol.136 pp.351-384
- Stroud, S. (1998) 'Moral Overridingness and Moral Theory' in *Pacific Philosophical Quarterly* 79
- Urmson, J. (1958) 'Saints and Heroes' in A. Melden (ed.) *Essays in Moral Philosophy*, Seattle: University of Washington Press
- Wallace, R. J. (2006) *Normativity and the Will: Selected Essays on Moral Psychology and Practical Reason*, Oxford: Oxford University Press
- Williams, B. (2011) *Ethics and the Limits of Philosophy*, Abingdon: Routledge.
- Woollard, F. (2010) Review: Peter Singer Under Fire, *Philosophical Reviews*

Available online at:

<http://ndpr.nd.edu/news/24252-peter-singer-under-fire-the-moral-iconoclast-faces-his-critics/>

Wolf, S. (1982) 'Moral Saints' *The Journal of Philosophy* vol.79 no.9 pp.419-439

Zimmerman, M. (2008) *Living With Uncertainty*, Cambridge: Cambridge University Press