

'Ought', 'Can', and Fairness

Rob van Someren Greve • Forthcoming in *Ethical Theory and Moral Practice*

(Final version available at <http://link.springer.com/article/10.1007/s10677-014-9492-1>)

Abstract According to the principle that 'ought' implies 'can', it is never the case that you ought to do something you cannot do. While many accept this principle in some form, it also has its share of critics, and thus it seems desirable if an argument can be offered in its support. The aim of this paper is to examine a particular way in which the principle has been defended, namely, by appeal to considerations of fairness. In a nutshell, the idea (due to David Copp) is that moral requirements we cannot comply with would be unfair, and there cannot be unfair moral requirements. I discuss several ways of spelling out the argument, and argue that all are unsatisfactory for a variety of reasons.

Introduction

According to the principle that 'ought' implies 'can', it is never the case that you ought to do something that you cannot do. Following a common way of making this idea more precise, let us take 'ought' to express *overall moral obligation*, and 'implies' to express *strict implication*, or entailment. The principle, then, reads as follows:

'Ought' Implies 'Can' (OIC): For any agent *S*, and action *A*, necessarily, if *S* is morally obligated to do *A*, then *S* can do *A*

Understood in this way (modulo some modifications that need not concern us here¹), many philosophers accept OIC for at least *some* sense of 'can'.² Just what sense of 'can' is implied by 'ought' is a matter of debate: some take it to be a sense that merely requires that it is nomologically possible that an agent perform an action that is morally required (or weaker still, that this is merely logically or metaphysically possible),³ while others endorse stronger views, requiring for example that the agent *knows how* to perform an action if that action is

¹ In order to deal with cases of self-imposed inability, some add two time-indices, so that the principle, when fully spelled out, reads as follows: 'for any agent *S*, action *A*, and times t_i , t_j , if *S* at t_i has an all-things-considered moral obligation to perform *A* at t_j , then at t_i it is true that *S* can do *A* at t_j . For this point, see Zimmerman (1996: 95-113), Howard-Snyder (2006: 235), and Vranas (2007: 175-78).

² There are various other ways to interpret the principle; Vranas (2007) offers an extensive discussion and numerous references. Note that Vranas himself defends the principle for *pro tanto* as opposed to overall moral obligation, but this is a minority view.

³ For the requirement that it is merely nomologically possible that an agent perform a required act (if the act indeed is morally required) see Streumer (2007).

morally required, or requiring that the agent can perform a required action *intentionally*.⁴ How exactly to interpret OIC is not my present concern; rather, the aim of this paper is to examine a particular way of arguing for OIC.

Although many accept the principle, it also has its fair share of critics,⁵ and thus it is worth examining whether a defense can be mounted on behalf of the principle. The defense of OIC that I want to discuss here derives from the work of David Copp (2003, 2008); we can call it ‘the argument from fairness’ for OIC. In outline, Copp’s argument is that moral requirements we cannot comply with would be unfair, and there cannot be unfair moral requirements, hence, there are no moral requirements we cannot comply with. I discuss two ways of spelling this argument out, and argue that both fail (Sections 1 and 2). After that, I present a more general problem for appeals to considerations of fairness in arguing for OIC in Section 3: these arguments, I claim, inherently show too much. Section 4 considers whether there is a way of formulating the argument from fairness that does not run into the objection from Section 3, and argues there is not.

1. The Argument From Fairness: First Version

To start, consider how David Copp summarizes the argument from fairness for OIC:

First, an adequate moral theory would imply ... the proposition that agent-requirements are morally unfair if the person required to act in a certain way is unable to act in that way. Second, a moral theory would be incoherent if it both implied this proposition and failed to rule out the proposition that a person can be morally required to act in a certain way even if she is unable to act in that way. Hence an adequate moral theory must reject the latter proposition, which means, in effect, that an adequate theory must imply [OIC] (2003: 272, cf. 2008: 71).

Some clarifications are in order. First, by ‘agent-requirements’, Copp means expectations that one agent has regarding the conduct of another; as an example we are invited to think of ‘a situation in which a boss requires an employee to do something that the employee lacks the ability to do’, for instance if ‘[a] supervisor at the post office’ demands ‘that a mail carrier cook a soufflé for everyone in the post office in the next five minutes when the employee does not even know what a soufflé is’ (Copp 2003: 271). It would be unfair of the boss to demand that the employee do this, and expect that the demand is met; the suggestion is that the correct moral theory must imply that this would be an unfair demand, and that it must have this implication because it must contain or imply a general principle which states that it is unfair for one agent to demand that another agent do something that the second agent cannot do, in whatever sense of ‘can’ we think is relevant here.

Second, the notion of fairness that figures in this claim is not a distributive notion: even if the boss of the post office indiscriminately requires of all of his employees that they

⁴ For the view that ‘ought’ implies ‘knows how’, see Bergström (1996) and Howard-Snyder (1997); for the view that ‘ought’ implies ‘can do intentionally’, see Wiland (2007).

⁵ Notable denials of OIC on any reading close to the one offered in the text include Stocker (1971), Sinnott-Armstrong (1984), Saka (2000), and Graham (2011).

do things they cannot do, the particular employee who is expected to make a soufflé in the next five minutes can still object to this particular demand on grounds of unfairness (cf. Graham 2011: 365). Let us grant, for the sake of argument, that there is a non-distributive notion of unfairness,⁶ and that this notion can be appealed to in explaining why there is something inappropriate about or amiss with agent-requirements pertaining to actions that those upon whom an agent-requirement is imposed cannot perform, in whatever sense of ‘can’ we think is relevant here.

How does all of this bear on whether OIC is true? Copp’s idea seems to be this. The correct moral theory, he suggests, must imply the general principle just noted, which we can call ‘Fair Demands 1’, or ‘FD1’ for ease of reference, and summarily state as follows:

Fair Demands 1 (FD1): For any agents S_1 , S_2 , and action A , it is unfair for S_1 to demand, require, or expect that S_2 do A if S_2 cannot do A

Copp holds that if the true moral theory implies FD1, it must, on the pain of incoherence, also imply the contrapositive of OIC:

‘Cannot’ implies ‘not ought’: For any agent S , and action A , if S cannot do A , then it is not the case that S is morally obligated to do A

And from this claim, OIC follows. Why, though, is it incoherent to accept FD1 and reject OIC? Neither claim obviously entails the other; they appear to be logically independent. If there is a problem of incoherence lurking somewhere, we must do a little work to uncover it. Here is one way to spell the argument out. Assume, for *reductio*, the following:

1. Morality demands, of some agent S , that S do A , and S cannot do A

Claim (1) is equivalent to ‘there is an agent S and action A such that S is under an obligation to do A and S cannot do A ’, and thus it is equivalent to the denial of OIC. If ‘morality’ can occupy the place of an agent-variable in FD1, we can infer from that claim and (1) that, therefore,

2. There is an agent S such that morality makes unfair demands of S

Finally, we can add the following, reasonable-sounding premise:

⁶ This assumption is not uncontroversial. To illustrate why it is not, consider the following famous anecdote:

During a trial about alleged police brutality, a lawyer asked Sydney [Morgenbesser] under oath whether the police had beat him up unfairly and unjustly. He replied that the police had assaulted him unjustly, but not unfairly. The lawyer was puzzled. “How is that possible?” he queried. “Well,” Sydney reportedly said, “They beat me up unjustly, but since they did the same thing to everyone else, it was not unfair.”

(This report is taken from Fletcher [2005: 548].) What Morgenbesser seems to suggest is that fairness is *essentially* distributive, i.e., that it makes no sense to say of some agent S that S has been treated unfairly if all of the other relevant agents received a similar treatment as S received, although S can be treated unjustly under such circumstances.

3. It is impossible that morality makes unfair demands of any agents

As Copp puts it, 'it is not intelligible to suppose that morality itself is morally unfair' (2003: 274). The above represents a way of putting that thought to use in a reductio of the rejection of OIC: if the demand that *S* do *A* is unfair if *S* cannot do *A*, and if morality cannot make unfair demands, then morality cannot demand that *S* do *A*, contrary to what (1) asserts. If, on the other hand, morality does demand that *S* do *A*, then either this demand is not unfair (and thus FD1 is false), or morality can make unfair demands (and thus (3) is false). Since both FD1 and (3) are intuitively plausible, it seems that (1) must be rejected.

In its present form, however, the argument from fairness for OIC suffers from a fatal flaw. Recall that, in getting from (1) to (2), we had to assume that morality can occupy the place of an agent-variable in FD1, and that assumption surely looks like a category mistake. As Nomy Arpaly formulates the problem in commenting on Copp's argument, this 'is where Anscombe would probably suspect us of being under the spell of regarding morality as a set of commands from a celestial boss, who can be fair or unfair to us' (2006: 107). It seems safe to assume that there is no 'celestial boss' whose orders constitute or imply what we, earthly mortals are morally required to do; given this assumption, we cannot personify morality and rely on a principle governing the fairness of demands made by *agents*. And that means that we cannot infer (2) from (1) and FD1. If considerations of fairness provide support for OIC, this first version of the argument from fairness does not offer a plausible way of spelling out how they manage to do so.

2. The Argument From Fairness: Second Version

Copp himself is not unaware of the problem facing the argument from fairness in its first version, as he notes that '[w]hen a person is *morally* required to do something there is not normally some *agent* who is requiring her to do it' (2003: 271, emphases in original). But as Copp sees things, this hardly means we cannot bring the adduced observations to bear on the question of whether OIC is true. There is a different, less direct way of doing so than that sketched in the previous section:

[I]f there would be no unfairness in the fact that a person is morally required to do something she is unable to do, then it is unclear why it would be unfair, or where the unfairness would be, in an agent's requiring a person to do something who is unable to do it (Copp 2003: 272).

One way to read this passage is to take Copp as suggesting that we ought to reverse the order of explanation: what is primary is the idea that a demand or requirement (etc.) is *as such* unfair whenever the action that is required is something that an agent cannot do, and that it is *this* principle which explains why it is unfair if an agent requires or expects another agent to do something that the latter agent cannot do.⁷ We can construct the argument from

⁷ I am not sure Copp would agree to this way of interpreting the quoted passage, for he does not distinguish different ways of understanding the argument from fairness. Be that as it may, it is worth considering this

fairness for OIC as an inference to the best explanation: FD1 is true, and something that stands in need of explanation; the best explanation of why FD1 is true is that the following principle, 'Fair Demands 2', is true as well:

Fair Demands 2 (FD2): For any agent *S*, and action *A*, if *S* is required to do *A*, and *S* cannot do *A*, then the requirement that *S* do *A* is unfair

The second version of the argument from fairness now goes as follows: since FD3 figures in the best explanation of a true claim (namely, FD1), FD2 is or can at least be assumed to be true, and since morality cannot be unfair (as claim (3) states), there cannot be obligations to do something that one cannot do, and so therefore, OIC is true. This version of the argument avoids personifying morality, for FD2 is not a principle about agent-requirements, but instead one that covers requirements in general (including moral requirements). In this respect, then, the third version is a clear improvement over the original argument.

While an improvement over the first version, the second version of the argument is hardly problem-free. For one, we can object that our intuitive grasp of the notion of an 'unfair demand' stems from considering the fairness of agent-requirements. We can make sense of this notion in the context of one agent requiring another to do something – for instance, when the action required is one that the agent cannot perform. While we can perhaps abstract from our grasp of this type of case to the more general idea of an unfair requirement figuring in FD2, i.e., a requirement that is not issued by an agent, if we were to reverse the order of explanation, we would get things backwards conceptually. It is no accident that Copp starts out with an example of an unfair demand made by an agent on another agent in introducing the argument from fairness, as thoughts about the unfairness of requirements have their home in the context of thinking about agent-requirements. And this may lead us to question whether FD2 really is a good explanation of FD1.

Admittedly, by itself this first point is merely suggestive, as the order of discovery need not be identical to the order of justification. It can still be true that what accounts for unfairness in the case of agent-requirements is the fact that any requirement is as such unfair if it concerns an action that the agent who is under the requirement cannot perform. We can make the point more robust if we can defend the claim that it does not really make sense to assess moral requirements as fair or unfair. Here is why I find this idea attractive. We can and do use the same term, 'requirement', both for describing an expectation that one agent has about the behavior of another, and for the deontic status of an action (recall that we assumed that 'doing *A* is obligatory for *S*' is equivalent to '*S* is morally required to do *A*'). However, the properties of actions that we refer to in these two different uses of the term do not fall into the same ontological category. Used in the first way, the feature referred to is a relation between facts about the mental state that some particular agent is in, viz., the fact that the agent who issues the requirement expects another agent to act in a particular way, and a particular action; the other is a moral feature of an action. These are distinct sorts of properties, because instantiating the latter property entails that there are good reasons to act in accordance with the requirement, whereas instantiating the former does not – although there are of course situations in which the fact that one agent expects

version of the argument for its own sake, as it avoids the issues with the previous version. A different reading is considered in Section 4.

another to act in a certain way, when this is combined with certain other facts, entails that there are good reasons to act in the way expected. We might put this by saying that while the property referred to by a moral requirement is *normative*, the property referred to by an agent-requirement is not; this surely marks an important ontological divide.

Since the term ‘requirement’ refers to radically different features that an action may have in these different uses of it, and if I am right that our understanding of the charge of unfairness has its home, so to speak, in the context of thinking about agent-requirements, then why is the charge, which we are familiar with when it pertains to certain mental states of agents, also one that can apply to properties that actions may have which fall into a fundamentally different ontological category? If there is an answer to this question, I do not know what it is supposed to be. If there is no answer, then FD2 is *not* a good explanation of FD1: FD1 attributes a moral property (‘unfairness’) to a non-moral property (viz., the relation between an action and the mental state of an agent), whereas FD2, provided it is assumed to bear on the question of whether OIC is true, attributes a moral property (‘unfairness’) to a further moral property (‘obligatoriness’). We might put this point by saying that FD1 and FD2 are principles with fundamentally different subject-matters, and thus the latter principle cannot help explain why the former principle is true, if indeed that principle is true.

If this observation is on target, then the main source of support for FD2 – that is, its figuring in the best explanation of FD1 – drops out, and with that, the third version of the argument from fairness falters. Furthermore, if FD2 indeed lacks support, we presumably also cannot appeal to this principle in trying to rescue the second version of the argument.

3. A General Objection

Drawing on the points raised in the previous section, it is worth noting that there may be a more general problem with arguments for OIC that appeal to considerations of fairness. To start, we could object that, like the first, the second version of the argument involves a category mistake as well: what can be fair or unfair is what we expect of others, how we treat them, how we distribute goods among them, and so on, but ascribing fairness or unfairness to *deontic facts* (i.e., an action’s being obligatory, or wrong) looks simply incoherent – deontic facts or properties are just not the sort of thing that can be fair or unfair. While Copp is correct when he writes that ‘it is not intelligible to think that morality itself might be morally unfair’, the proper conclusion to draw from this observation is not that morality itself therefore has to be fair, but rather that assessments of fairness and unfairness make no sense when these concern the moral status of acts. If this is correct, FD3 must be rejected as meaningless, and therefore as incapable of supporting any conclusion whatsoever.⁸

⁸ Shelly Kagan offers a comparable diagnosis of the suggestion that ‘morality’ might violate a constraint against harming if there are no options to pursue one’s own interest; as he puts it, ‘here we have a case in which the personification of morality can lead us into error’ (1989: 208). I am inclined to maintain that the same can be said if we formulate the argument from fairness such that it concerns the *demands* that morality makes, as opposed to putting it in terms of what morality requires: were we to claim that the demands that morality makes are unfair, we are again misled by our way of speaking – we are applying a criticism that is sensible when applied to demands made by a particular agent, but not to deontic facts, such as an act’s being

This is suggestive, but hardly a knock-down objection. A more compelling objection is this. Both versions of the argument from fairness are of the form ‘it would be unfair if *X*, therefore, not-*X*’ (which is made explicit by claim (3) above). Clearly, we cannot just put anything in place of *X* in this schema and end up with a valid argument – it would be unfair, for example, to pay men and women different salaries for performing the same tasks, but from that fact we cannot infer that therefore men and women are not rewarded unequally for performing the same task. But what of instances where ‘*X*’ is a moral claim? Elsewhere, I have argued against the view that arguments of the form ‘it would be good if *X*, therefore *X*’ are valid when *X* is a evaluative or moral claim.⁹ In a nutshell, the problem is that such arguments would prove too much: assuming that we do not have the means to eradicate poverty, it would be good if poverty (and all associated evils) were not bad, but good – after all, the world would be a much better place if there were less bad things in it. If we can move from ‘it would be good if *X*’ to ‘*X*’ when *X* is a claim of the form ‘*Y* is good’, we could infer that widespread poverty is not bad, but good. But that conclusion is absurd, and what needs to go in order to be able to avoid being saddled with it is the view that arguments of the form ‘it would be good if *X*, therefore *X*’ are valid when *X* is a evaluative or moral claim.

Effectively the same problem arises for the view that arguments of the form ‘it would be unfair if *X*, therefore, not-*X*’ are valid when *X* is a moral claim, e.g., a claim of the form ‘*S* is obligated to do *A*’: such arguments prove too much. Say that Sally is a member of the team of lifeguards at the community pool, and has been doing more than her fair share of saving neighborhood children from drowning. Jimmy and Billy, the other members of the team, have by contrast not done their part all summer, and today, they will once again just sit around and let Sally do all the work. A child is in danger of drowning – is Sally obligated to rescue it? Since Jimmy and Billy will not lift a finger, and Sally knows that this is so, it seems she is, even though it is, in a sense, clearly unfair that she is, given that she has been the only one to do her part so far. Yet from that we cannot conclude that therefore she is not under an obligation to save child. And the point generalizes, I believe: it is often unfair that the burden of doing something falls on a specific person’s shoulders, solely because others who could have done their part did not in fact do so. Unfair as that may be, it seems unwarranted to conclude that therefore these people are not under an obligation to act, provided there is enough at stake. Were we to accept the suggested argument-schema as valid for instances where *X* is a moral claim, we would have to conclude that Sally is not obligated to save the child, and likewise in other situations like it. But that, I submit, just goes to show we should not accept this schema as valid even when *X* is a moral claim.

If this objection succeeds, any variation of the argument that requires claim (3) fails, as any such argument appeals to the idea that there cannot be unfair moral requirements. If there can be such requirements, it would be a moot point to show that a convincing case can be made that a moral requirement to do what one cannot do would be unfair. The

morally obligatory. While we can express the fact that a certain act is required by saying that ‘morality demands that *S* performs this act’, we should not take the fact that this is a legitimate way of speaking lead us into thinking that the sort of criticism that applies to ‘personal demands’ also applies to morality’s demands. (Thanks to an anonymous referee for pressing me on this point.) Note, though, that if I am wrong, and we can meaningfully criticize morality’s demands as unfair, an argument building on this sort of criticism will run into the second, more compelling objection formulated in the text.

⁹ For defenses of this view, see Nagel (1995: 91-92) and Enoch (2009). For a discussion of the problems facing this view, see my (2011).

argument from fairness fails, then, even when we grant that the other objections I have raised do not succeed.

4. An Alternative Argument From Fairness?

Does the argument from fairness need the premise that morality cannot be unfair? Take again the passage cited at the beginning of Section 2:

[I]f there would be no unfairness in the fact that a person is morally required to do something she is unable to do, then it is unclear why it would be unfair, or where the unfairness would be, in an agent's requiring a person to do something who is unable to do it (ibid., 272).

There is a different, and perhaps more straightforward way of explicating the argument suggested here.¹⁰ Suppose that the following slight reformulation of claim (1) is true:

1*. There is some agent, S_1 , who is morally required to perform an action, A , and S_1 cannot perform A .

Suppose in addition that a second agent, S_2 , requires that S_1 performs A (in the sense that S_2 expects S_1 to perform A , and S_2 expresses this expectation to S_1). Is it unfair for S_2 to expect S_1 to perform A under these conditions? Recall FD1:

Fair Demands 1 (FD1): For any agents S_1 , S_2 , and action A , it is unfair for S_1 to demand, require, or expect that S_2 do A if S_2 cannot do A

If we accept FD1, we have to answer this question positively. Now, if we accept both (1*) and FD1, it appears we have to reject the following principle, 'Fair Demands 3':

Fair Demands 3 (FD3): For any agents S_1 , S_2 , and action A , it is not unfair for S_1 to demand, require, or expect that S_2 do A if S_2 is morally required to do A

If we think FD3 is true, we can offer the following alternative argument from fairness: it is unfair for S_1 to require that S_2 do A when S_2 cannot do A (FD1), and it is never unfair for S_1 to require that S_2 do A when S_2 is morally required to do A (FD3), therefore, S_2 is never morally required to do A when S_2 cannot do A (not-1*). This alternative argument does not rest on the idea that there cannot be unfair moral requirements; its premises are compatible with the claim that there are such requirements, and thus it does not run into the problem that any variation on the first and second version of the argument faces.

The alternative argument, however, has its own problems. Most immediate is the problem that FD3 seems clearly false: if S_1 has failed to fulfill his own moral obligations, it would be hypocritical, and thus presumably unfair, of S_1 to require that S_2 fulfill *her* obligations (would it be fair for Billy or Jimmy to require that Sally rescue the drowning

¹⁰ Thanks to an anonymous referee for suggesting this alternative way of spelling out Copp's argument.

child?). There are, that is, clearly cases in which it is intuitively unfair for one agent to require that another agent do what she is morally required to do, where we have no doubt about the existence of the underlying moral obligation.

In response, one could amend FD3 so as to accommodate this type of case:

*Fair Demands 3** (FD3*): For any agents S_1 , S_2 , and action A , it is not unfair for S_1 to demand, require, or expect that S_2 do A if S_2 is morally required to do A , provided S_1 has not failed to fulfill his or her own moral obligations

This revised principle suffices to get the argument from fairness of the ground, as there is no reason to think that all cases in which one agent is barred from requiring of another that she do what she is morally required to do are cases in which the first agent has failed to fulfill his own obligations. In other words, from FD1 and FD3*, we can still infer (not-1*).

The best response to this alternative argument is, I think, straightforward, although probably a little disappointing: if one accepts that there can be situations where someone is morally required to do something she cannot do, one should take those cases to provide a counterexample to FD1, and maintain that it *is* fair to require that someone does something even if she cannot do it, provided that she is morally required to perform the action in question. That is, if there are successful counterexamples to OIC, then there are successful counterexamples to FD1, for these are one and the same. If, by contrast, there are no successful counterexamples to OIC, then there are no successful counterexamples to FD1 either. To illustrate: if cases of self-imposed inability, where an agent makes it impossible for herself to, say, keep her promise to return a book she borrowed by burning the book, show that OIC is not true (because the agent is still under an obligation to return the book), then these cases presumably also show that it is not unfair to require that an agent do what she cannot do. Assuming the agent is still obligated to return the book despite having burned it, it does not seem unfair for its owner to require that she return it. If such cases do not provide counterexamples to OIC (because the agent is no longer obligated to return the book when it has become impossible for her to do so), they also do not show that we can fairly require that someone do something she cannot do.

It has not been my aim in this paper to arbitrate whether there are successful counterexamples to OIC; what I set out to do was evaluate whether appealing to considerations of fairness can figure in a successful argument for OIC. Even without taking a stand on the success of purported counterexamples to OIC, the observations offered here suggest that it will just not do to appeal to such considerations in the way sketched in this section, for what considering FD3* brings out is that we cannot but assess the intuitive plausibility of FD1 *simultaneously* with that of OIC. Whatever the intuitive plausibility of FD1, this is not a claim that can be put to work in defending OIC against its critics, for they should (and presumably will) treat cases that show OIC to be incorrect to show FD1 to be false as well. For this reason, the alternative argument from fairness fails, just like the first two versions considered above. In addition, these observations suggest another objection to the first two versions of the argument from fairness, as both rely on FD1 as well, thus further undermining the prospects of a plausible appeal to considerations of fairness in a defense of OIC.

5. Conclusion

In this paper, I set out to evaluate whether there is a plausible way of appealing to considerations of fairness in defending the principle that 'ought' implies 'can' (OIC). I have argued there is not: construed in one way, such an appeal runs the risk of personifying morality, and even if we avoid this pitfall, it mistakenly attributes a moral property to a further moral property. In addition, the appeal requires that arguments of the form 'it would be unfair if X , therefore, not- X ' are valid at least whenever X is a moral claim, but we should reject this argument schema as invalid across the board. An alternative construal of the appeal to fairness does not fall prey to these objections, but considering it brought out that a crucial premise of the argument (which, incidentally, is also a crucial premise in the first two versions of the argument) is not independent of OIC, as its intuitive plausibility is arguably directly tied to that of OIC: if there are successful counterexamples to OIC, there are successful counterexamples to the principle that one cannot fairly require that another do something she cannot do. Whether there are such examples is a subject for a different occasion, though.

Acknowledgements

Thanks to Emily Given for discussion, and two anonymous referees for this journal for their helpful comments on a previous version of this paper.

References

- Arpaly, N. (2006). *Merit, Meaning, and Human Bondage: An Essay on Free Will*. Princeton, NJ: Princeton University Press.
- Bergström, L. (1996). 'Reflections on Consequentialism'. *Theoria* 62 (1-2), pp. 74-94.
- Copp, D. (2003). "'Ought' Implies 'Can', Blameworthiness, and the Principle of Alternate Possibilities'. In McKenna, M. and D. Widerker (eds.), *Moral Responsibility and Alternative Possibilities* (pp. 265-299). Aldershot: Ashgate Press.
- Copp, D. (2008). "'Ought' Implies 'Can' and the Derivation of the Principle of Alternate Possibilities'. *Analysis* 68 (1), pp. 67-75.
- Enoch, D. (2009). 'Wouldn't it be Nice if p , Therefore, p (for a Moral p)'. *Utilitas* 21 (2), pp. 222-224.
- Fletcher, G. P. (2005). 'Justice and Fairness in the Protection of Crime Victims'. *Lewis & Clark Law Review* 9 (3), pp. 547-557.
- Graham, P. (2011). "'Ought' and Ability'. *Philosophical Review* 120 (3), pp. 337-382.
- Howard-Snyder, F. (1997). 'The Rejection of Objective Consequentialism'. *Utilitas* 9 (2), pp. 241-248.
- Howard-Snyder, F. (2006). "'Cannot' Implies 'Not Ought'". *Philosophical Studies* 130 (2), pp. 233-246.
- Kagan, S. (1989). *The Limits of Morality*. Oxford: Oxford University Press.
- Nagel, T. (1995). 'Personal Rights and Public Space'. *Philosophy and Public Affairs* 24 (2), pp. 83-107.

- Saka, P. (2000). 'Ought Does Not Imply Can'. *American Philosophical Quarterly* 37 (2), pp. 93-105.
- Sinnott-Armstrong, W. (1984). "'Ought' Conversationally Implies 'Can'". *Philosophical Review* 93 (2), pp. 249-261.
- Stocker, M. (1971). "'Ought' and 'Can'". *Australasian Journal of Philosophy* 49 (3), pp. 303-316.
- Streumer, B. (2007). 'Reasons and Impossibility'. *Philosophical Studies* 136 (3), pp. 351-384.
- Van Someren Greve, R. (2011). 'Wishful Thinking in Moral Theorizing: Comment on Enoch'. *Utilitas* 23 (4), pp. 447-450.
- Vranas, P. B. M. (2007). 'I Ought, Therefore I Can'. *Philosophical Studies* 136 (2), pp. 167-216.
- Wiland, E. (2007). 'How Indirect Can Indirect Utilitarianism Be?' *Philosophy and Phenomenological Research* 74 (2), pp. 275-301.
- Zimmerman, M. J. (1996). *The Concept of Moral Obligation*. Cambridge: Cambridge University Press.