

JUSTIFYING THE OPTION TO ACT SUB-OPTIMALLY¹

1. INTRODUCTION

Morally speaking, I probably shouldn't be writing this essay. I could be doing something much better with my time. I'm not skilled, but I'm a fast learner—perhaps with dedication I could save some lives. I could surely earn more than I do now, and money can *definitely* save lives. We can quibble over details, but it would be a remarkable coincidence if sitting here at this desk, in my warm study, with the rain beating its steady rhythm on the roof of my house, I was doing the very best I could, morally speaking.² What luck that would be!

The thing is, I don't feel like I'm doing anything wrong. Not really at all. I should probably give more than I do to charity. I definitely spend more than I should on silly things like Spider-Man and Captain America figures (for my son... mostly). But I'm not going *that* far wrong.

To give it a technical name, I think that I have an *agent-favouring option to act sub-optimally*, which licenses my indulging my philosophical curiosity rather than realising the morally best outcome. Of course, if I want to sacrifice my own interests for the sake of others, then I would be permitted to do so. If I did so, I would be acting *supererogatorily*.³

Agent-favouring options are not the only kind.⁴ Suppose, for example, that you and I have been slogging through the desert for hours, and my greater size means that I am more dehydrated than you. But I give you the last of our water. This too is a suboptimal outcome—the world would go better if I took the water rather than you. But altruistic self-sacrifices like this are clearly permitted.

Indeed, even *mere* self-sacrifice is morally permissible. Suppose I can order Chinese or Indian takeaway tonight. I would enjoy the Indian meal more, and my innocent enjoyment makes the world a better place. So if I was required to maximise value, it would be morally wrong to order the Chinese. But this is absurd.

I would go even further: I can harm or frustrate my own interests, without acting morally impermissibly. If nobody else is affected by my

¹ Seth Lazar, ANU. This is a first-and-a-bit draft, so please do not circulate. Comments most welcome, to sethlazar@gmail.com. Thanks to Doug Portmore for comments on the first-first draft.

² The quibblers are the advocates of various forms of indirect or subjective consequentialism, who argue that actually just living our lives in a more or less ordinary way *does* maximise value. See, for example, Railton [1984]; Jackson [1991]; Pettit [1997].

³ Urmson [1958] is the locus classicus for discussion of supererogation.

⁴ For arguments that commonsense morality includes agent-sacrificing options (though not arguments defending that inclusion), see Broad [1971]; Stocker [1976: 207]; Slote [1984a: 180]; Sider [1993]; Gert [2000: 235].

self-sacrifice, then I think it is nobody else's business—not even morality's.⁵ I'm morally free to act totally irrationally, as long as my doing so affects only myself. Some will think that the scope of these *agent-sacrificing* options is more limited than this (my views on this are perhaps extreme). But all should agree that we have *some* such options.

Let's call all of these simply *moral options*. Now, perhaps my readiness to assent to moral options is mere motivated reasoning. Or perhaps it is an illicit importation of a certain kind of political liberalism into moral theory. Maybe it's all self-serving nonsense; really I ought to maximise value—I don't have options to act sub-optimally. We should not dismiss these possibilities. But before we embrace them, there are at least three philosophically interesting questions we can ask about moral options.

First, does commonsense morality recognise moral options?⁶

Second, are moral options defensible?⁷

And third, can consequentialists—people who normally think that morality requires us to maximise value—accommodate moral options?⁸

In my view, the first and third of these questions have been amply covered. But the second has been neglected. More, perhaps, has been said about agent-favouring options than about agent-sacrificing ones. But their justification has been explored in most detail by someone convinced of their immorality.⁹ Though his book, *The Limits of Morality*, inspired several defences of moral options against his critique,¹⁰ I think Kagan [1994: 928] rightly ended those discussions by noting that, as yet, nobody had given a persuasive defence of moral options that coheres with a plausible overall moral theory. The two decades since have not seen significant advances.¹¹ Most philosophers working in this area have simply tried to develop a plausible extension of consequentialism that could accommodate moral options, if they were justified. This is understandable, since moral options are at least *prima facie* an embarrassment for consequentialism. But it means that we lack a convincing positive case in their favour.

In this paper, I want to provide just that: a new argument for moral options.¹² The basic idea is very simple, and has already been alluded to. My

⁵ Indeed, even if others are affected, but only by my frustration of their me-regarding interests, I think I have a total licence to do what I like.

⁶ See especially Urmson [1958]; Broad [1971]; Stocker [1976]; Slote [1984a]; Kagan [1989]; Scheffler [1994]; Cullity [2004]; Mulgan [2005]; Hurka and Shubert [2012].

⁷ See, especially, Scheffler [1992, 1994]; Hurley [1995].

⁸ See especially Sider [1993]; Cullity [2004]; Mulgan [2005]; Sobel [2007]; Rogers [2010]; Vessel [2010]; Portmore [2011]; Dorsey [2012, 2013]; Suikkanen [2014].

⁹ Kagan [1989].

¹⁰ Brock [1991]; Kamm [1992]; Bratman [1994]; Waldron [1994].

¹¹ Though see Hurley [1995].

¹² The argument has affinities with ideas sketched by Quinn [1989a]; Kamm [1992]. It is close in spirit, if not in its details, to Hurley [1995]. Fiona Woollard [2013] uses a similar idea in defending the doctrine of doing and allowing, though she argues that we have ownership rights over our bodies, which I do not endorse. Victor Tadros [2011] also uses

life, and in particular my self-regarding interests, are my own. I get to have more of a say over how they are disbursed than does anyone else. If I want to sacrifice my own interests, even though doing so makes an outcome overall worse, then (within bounds) I'm entitled to do so, *because they're my interests*. If I want to preserve my own interests, even though doing so misses an opportunity to make the outcome better than, within bounds, I'm entitled to do so, *because they're my interests*. I have a special authority over my own interests. This authority is in part constitutive of moral status. It is (part of) what protects me against being a mere site or tool for the realisation of value.

In section 2, I'll set out this argument in more detail. I'll also show how it differs from existing arguments for moral options, and how it evades Kagan's main objections to those views. In section 3, I'll formulate a principle that operationalises the abstract ideas set out in section 2. Although it remains quite abstract, it at least provides a method for determining which moral options to endorse. In section 4, I'll compare my approach with the main competitors. These have been advanced by consequentialists seeking to accommodate options within their theory, without doing too much violence to its basic commitments. I'll argue that my approach better fits the intuitive data, has more explanatory power, and is simpler than the alternatives. The last section will consider objections.

2. NOT MERE SITES FOR THE REALISATION OF VALUE

The first stage of my argument for moral options identifies a general set of moral considerations that are sufficient to justify options. The second renders this general view more specific. In doing so it incurs more philosophical commitments, but it gains precision and explanatory power. Having set these ideas out, I'll explain how they link to arguments considered by Kagan in his magisterial evisceration of moral options (I think it is different from them all). And I'll conclude the section by showing how my account vindicates both kinds of moral options, while avoiding Kagan's main objections.

My starting point is familiar. Persons have moral status. I won't explore in detail what grounds personhood, and why those grounds generate moral status. I have nothing much to add to familiar, broadly Kantian accounts, according to which our capacity for rational thought and decision, and perhaps our very ability to act morally, are what make persons special, what gives them status. Others may disagree about this. But whatever your views on what *grounds* moral status, I think you should agree

Quinnian ideas in his discussions of duties to aid. I am not aware of any attempt to apply these ideas to the justification of options.

that having moral status entails at least the following two facts: first, if you have moral status, then your flourishing or suffering *matters*. Your interests can give other people moral reasons for action. They have *value*. But second, if you have moral status, then *you* matter, independently of whether you flourish or suffer. *You* can give reasons for action, which are not wholly reducible to the fact that flourishing or suffering is instantiated in your life. Your happiness matters—but *so do you*. You are not merely a site of value, whose moral significance is exhausted by how your happiness and suffering contributes to, or detracts from, the overall value of the world.¹³

This alone is enough for a preliminary case in favour of moral options. A moral theory without options would attend only to one aspect of our moral status—the fact that our interests matter. It would treat us as nothing more than bearers of interests, sites for the realisation of value. We would have to sacrifice our interests whenever doing so realised a somewhat greater good, and we would be prohibited from doing so unless it was optimific. Our decisions would be fully determined by the balance of value realised in our lives and others'. Except in cases where the balance of value is the same either way, we would have no real freedom of choice.

Moral options play a similar role in morality to that played by moral constraints.¹⁴ If we each mattered only insofar as our interests contribute to the value of the world, then it would be permissible to harm someone else just in case doing so realised somewhat more value than not harming him. The reasons he can give others would be wholly reducible to the value instantiated in his life. This is wrong—it neglects the fact that, while his interests matter, *he matters too*. He is not a mere site for the realisation of value.

I am not, however, arguing that morality would 'violate a constraint' if it viewed us as resources for the benefit of others Kagan [1989: 207].¹⁵ I am not first positing constraints, then arguing that morality, like people, can breach these constraints. Instead, I am arguing that the two dimensions of moral status together entail that we have moral options, and are protected by moral constraints. Options and constraints have a common source.

Of course, many consequentialists will balk at this idea.¹⁶ Many agree

¹³ Although the underlying inspiration for this idea is obviously Kantian, this version of it owes most to Warren Quinn and Frances Kamm's work. E.g. '[P]ersons are not mere means to the end of the best state of affairs, but ends-in-themselves, having a point even if they do not serve the best consequences' Kamm [1992: 358-9]. See also Quinn [1989a: 307-9].

¹⁴ This is not to say that moral options and constraints are wholly symmetrical, as some argue (e.g. Heyd [1982]; Hurley [1995]; Thomson [2008]; Tadros [2011]).

¹⁵ My argument is also different from Hurley's [Hurley [1995] argument that *we ourselves* violate a constraint when we treat ourselves as being obligated to accede to marginal interpersonal tradeoffs.

¹⁶ Not all. Portmore [2011: 4] in particular is deeply concerned about how some consequentialist theories implausibly license marginal interpersonal tradeoffs.

with Bentham, Sidgwick, Raz, Kagan and others in between and since, who argued that having moral status means *only* that one's interests are of moral concern, nothing else.¹⁷ On their account, if my sacrifice is counter-balanced by a gain to another, then even if the net gain is tiny, I can be required to make it. As Sidgwick put it, 'the good of any individual is of no more importance, from the point of view of the universe, than the good of any other', and 'it is my duty to aim at good generally, so far as I can bring it about, and not merely at any particular part of it'.¹⁸ The same reasoning would prohibit agent-sacrificing options. If I'm required to aim at the good in general, then I can no more permissibly frustrate my own good than I can harm others.

We are probably close to normative bedrock, and perhaps neither side can hope to convince the other of its mistake. Some people think that moral status consists solely in having one's interests included in the overall balance. Others think that it generates additional reasons besides those grounded in one's interests. Some think we are mere sites for the realisation of value. Others do not. This may be a fundamental fault-line in ethical theory.¹⁹ But perhaps some have not yet picked sides. And perhaps getting a little more specific will help convince those waverers. So, on to the second stage of the argument.

To not be a mere site for the realisation of value is to be recognised as an independent being, with one's own life to lead, and a fundamental sphere of freedom in which one is immune from the criticism or interference of others. This freedom is realised by our having *authority* over our self-regarding interests: having the licence to sacrifice them, or to withhold them, simply because we choose to do so, and even though it is suboptimal.

My flourishing and suffering, insofar as they are constituted by the fulfilment or frustration of self-regarding interests, surely concern and affect me more than they do anybody else—these are *self*-regarding interests! I am uniquely able to influence them. They are intimately and intricately connected to my identity. If I suffer some purely self-regarding harm, then of course that (*pro tanto*) reduces the net balance of goodness over badness in the world. But it does so only by making *me* worse off. Nobody else is affected (by hypothesis). If I undergo some sacrifice for the sake of someone else, then even if the overall balance of value is improved, the cost falls on me. I should get *some* degree of say over whether to make that sacrifice.

The basic idea here comes from Warren Quinn, albeit that he used it to motivate the distinction between doing and allowing harm. He argued

¹⁷ See especially Kagan [1989: 252, 354].

¹⁸ Sidgwick [1981: REF].

¹⁹ One might think it a fault-line between consequentialists and deontologists, but this would be a mistake. Some consequentialists, such as Portmore [2011: 4] also reject the idea of marginal interpersonal tradeoffs.

as follows:

A person is constituted by his body and mind. They are parts or aspects of him. For that very reason, it is fitting that he have primary say over what may be done to them—not because such an arrangement best promotes overall human welfare, but because any arrangement that denied him that say would be a grave indignity. In giving him this authority, morality recognizes his existence as an individual with ends of his own—an independent *being*. Since that is what he is, he deserves this recognition. Were morality to withhold it, were it to allow us to kill or injure him whenever that would be collectively best, it would picture him not as a being in his own right, but as a cell in the collective whole. (Quinn [1989a: 309])

Now, besides addressing these ideas to a different subject matter, my position is slightly different from Quinn's. I take no view, in this paper, on whether we have this kind of authority over our bodies and minds. Though the idea is plausible, authority over one's interests seems more fundamental. After all, the contrary view is not strictly that our *bodies* are cells in some organic collective whole, but that our *interests* are—one more input into an aggregate, to be sacrificed or advanced insofar as it makes the collective whole better overall. I think our authority runs out when we can benefit others at no cost to ourselves, which might still involve using our bodies, so would be ruled out on Quinn's principle.²⁰ I also think that some self-regarding interests are not obviously corporeal, and can constitute the kinds of costs that ground moral options.

Though I disagree with Quinn about what beings with moral status should have authority over, we agree that *having authority over something* is necessary for one not to be merely a cell in the collective whole, a site for the realisation of value. In his defence of the doing/allowing distinction, Quinn further developed this idea. If a person could be harmed just in case doing so realises a marginally greater benefit for others, then

he has a moral say about whether his body may be destroyed only if what he stands to lose is greater than what others stand to gain. But then surely he has no real say at all. For, in cases where his loss would be greater than the gain to others, the fact that he could not be killed would be sufficiently explained not by his authority in the matter but simply by the balance of overall costs. And if this is how it is in general—if we may rightly injure or kill him whenever

²⁰ Quinn's use of this idea to ground the doing/allowing distinction is developed in greater detail by Woollard [2013], who sticks to the idea that what matters is having authority over one's body.

others stand to gain more than he stands to lose—then surely his body (one might say his person) is not in any interesting moral sense *his*. It seems rather to belong to the human community, to be dealt with according to its best overall interests. (Quinn [1989a: 308])

Again, where Quinn speaks of the person's body, I think it makes more sense to speak of his self-regarding interests, and where Quinn is arguing for special constraints against doing harm, I am arguing for moral options. But the underlying idea is the same: if we don't have *some* authority over our own interests, then we are not really independent beings, but are instead just inputs into a collective algebra. But if we have this authority, then we must be entitled to withhold our aid from others, when the cost to us is almost as great as the benefit to them, and we must be allowed to thwart our own interests if we want to, regardless of how that affects the overall value of a situation.

Of course, though we are not mere sites for the realisation of value, nor are we solipsistic monads. After all, *other people* have moral status too—they matter, and so do their interests. What's more, though we are not *mere* sites for the realisation of value, our interests *do have value*. This means that our authority over our self-regarding interests is not boundless. Which means, in turn, that there are limits on our agent-favouring and agent-sacrificing options. If the good I can do for others is great enough, relative to the cost to me, then my authority over my own interests is overridden, and I *am* required to help them. And perhaps if the harm I will inflict on myself is great enough, then my authority can again be overridden, and I can be required not to sacrifice my own interests.²¹

In *The Limits of Morality*, Kagan identified two broad families of argument for moral options. 'Positive arguments' claim that moral options are necessary for people to live flourishing lives, in which they can commit themselves to people and projects independently of whether doing so realises the most good overall.²² 'Negative arguments' view moral options as concessions to an unfortunate fact of human nature: we are simply too selfish and unimaginative to adopt the Sidgwickian 'point of view of the universe'.²³ I think Kagan's objections to each of these arguments are right. Commonsense moral options apply even when one's valuable projects

²¹ My view remains that there are no self-regarding duties, and that if nobody else is affected, then it is nobody else's business. But I accept that this is controversial, and do not want to tie the discussion to a defence of that controversial view.

²² The locus classicus is Smart and Williams [1973]. See also Kagan [1989: 233, 334]. Brock [1991: 912] argued, in response to Kagan's book, that 'options are necessary to the realization of the ideal of persons as autonomous agents in constructing and pursuing their own conceptions of the good'.

²³ Kagan [1989: 261, 71]. The most prominent of these have been advanced by Scheffler [1992, 1994].

and commitments are not at stake. And the belief that morality should adapt to our motivations is a controversial metaethical stance, which we can reasonably reject. And we can now add to Kagan's objections: whatever their merits as justifications for agent-favouring options, neither the positive nor the negative argument can plausibly justify agent-*sacrificing* options.

My argument is different from these.²⁴ I agree, of course, that our having authority over our self-regarding interests is likely to lead to good results. Nobody else will care for those interests like we will. And I agree, too, that people probably are simply incapable of adopting a Sidgwickian motivation, regarding themselves as mere 'cells in a collective whole'. But true though these points may be, they are not what justifies our having this authority. Instead, it's a basic principle of justice that each person should have a sphere of freedom in at least in this area, where their self-regarding interests are primarily at stake. Incorporating moral options into our moral theory reflects respect for the moral status of persons as independent beings, with their own lives to lead. It's not a way to preserve something of value, or an unfortunate concession to our limited moral capacities.

Options are grounded, then, in two closely related, and I think compelling principles. We are not mere sites for the realisation of value, or cells in a collective whole. We are independent beings, with moral status. In virtue of this, we have a special authority over our self-regarding interests, which means we can choose to forebear from sacrifices that would be overall optimal, as well as choose to sacrifice our interests when doing so is sub-optimal. This zone of freedom is justified not by its good results, or our unfortunate incapacities, but by its being the innate right of every being that has moral status. So Kagan [1989: 238] is wrong to suggest that proponents of options have to 'explain why it is better... that a given reaction be permitted rather than forbidden or required'. They have to do no such thing. We have moral options because a basic principle of justice entitles us to a zone of freedom where our self-regarding interests are concerned. Betterness is beside the point.

Kagan's challenge was to come up with a substantive defence of moral options that coheres with a plausible overall moral theory. I think the present argument does that. The underlying ideas are the same as those that ground moral constraints, the distinction between doing and allowing, and (with a little further development) the distinction between harming someone as a means and harming them either as a side-effect, or in a manner that deprives no benefit from their presence that one would not have enjoyed in their absence.²⁵ It justifies both kinds of moral options, tying them to two deeply compelling ideas. And I think it answers Kagan's

²⁴ Though it is somewhat closer to Hurley [1995], who also sees options, like constraints, as being grounded in facts about moral status.

²⁵ Quinn [1989b].

two central objections to arguments for moral options.

The first is that they must not *over-generate* such options (e.g. Kagan [1989: 78]). In particular, while we may be permitted to favour our own interests rather than help others, it's much less plausible that we can have moral options to harm others. Consider a life and death scenario: clearly I'm permitted not to save your life if doing so would cost me mine (but of course I'm permitted to sacrifice my life for yours if I want to); equally clearly I'm *not* permitted to kill you (other things equal) if doing so is necessary to save my life.

My argument readily accommodates this distinction. Indeed, as already noted, my justification for options is closely connected to a compelling account of the difference between doing and allowing harm (Quinn [1989a]; Woollard [2013]). It is harder to justify doing harm than allowing it, because to actively harm someone in the course of pursuing one's own good would be to treat him as though he has no special say over his own interests, whereas allowing someone to suffer harm does not. I cannot defend that difference in detail here. But my special authority over my own interests would not license my thwarting the interests of others to advance my own good; quite the contrary.

Kagan's second concern (e.g. Kagan [1989: 253, 371ff, 1994: 338]) is that arguments that justify options to favour or sacrifice one's interests should not end up entailing requirements, rather than genuine options. He puts the basic worry like this: if what we ought to do is a function of the balance of our reasons for action, and if an argument shows that the balance of reasons favours, for example, preferring one's own lesser interests over another person's greater interests, then how do we avoid the conclusion that one is morally required, not merely permitted, to prefer one's interests? Kagan admits that some moral reasons could be 'noninsistent', so render something permissible, but not require it.²⁶ But though he recognises the conceptual space for such reasons, he denies that any of our moral reasons have this form (Kagan [1994: 381]).

I think my account resolves this problem. Our self-regarding interests generate noninsistent reasons because we have authority over them. If these reasons were insistent, so that I was morally required to prefer my own interests rather than perform the supererogatory act, or (bizarrely) required to harm myself when doing so is suboptimal, then I would not have any authority over my self-regarding interests. It would not be up to me to decide what to do. My self-regarding interests can justify action, but they cannot require it, because they are my interests, and I have authority over them! If I want to sacrifice myself, or to refrain from self-sacrifice, then within bounds I have the right to do so.

Recall Quinn's analogy with property rights. If I own this laptop, then

²⁶ The 'insistent/noninsistent' reason distinction is much the same as Gert's later distinction between the requiring and justifying strengths of normative reasons, e.g. Gert [2000, 2007].

I get a say over how it is used. The mere fact that it could do someone else more good than it does me is inadequate grounds for me to be required to give it to her. It's my laptop! Equally, if I fail to maintain it, or even smash it up, then I am probably entitled to do so, even though it is worse for us both. But obviously my licence to damage my own property does not mean that I am morally required to do so. And my right to hold onto it even if I could do more good by giving it away does not prohibit me from donating it if I want to. My authority over my laptop justifies my having agent-favouring options to keep it, and agent-sacrificing options to give it away or damage it. But it doesn't generate any requirements. The same is true for my authority over my self-regarding interests.

Of course, sometimes we do have agent-relative moral reasons to favour our own projects and commitments—reasons for which whether they apply, and the force that they have, depends on who is acting on them. But these are, I think, 'insistent' or 'requiring' reasons. My reasons to promote my loved ones' interests often ensure that doing so is genuinely the best I can do (taking agent-relative reasons into account) and is therefore required (provided it does not involve an excessive sacrifice of my self-regarding interests).

3. A CANDIDATE PRINCIPLE

The foregoing argument should show that any sensible moral theory should include *some* moral options. However, an important test of its success is whether it justifies *the right* options. If it could motivate only a subset of the plausible cases, or if it justified implausible options, then that would place it in doubt. So we need to go beyond the foundations of normative ethics, and consider a principle that puts these basic ideas into practice.²⁷ This will also help us contrast my proposal with principles put forward by others. I'll first present my principle, then explain what each of its elements means, before showing how it is derived from the argument above.

COST: An act is permissible if and only if either (a) there is no morally better act that has reasonable marginal costs to the agent or (b) it falls short of every such reasonable alternative only in virtue of costs borne by the agent.

I will not try to define every element in COST. I will not say, for example, what I mean by an act. And though I think we can naturally extend COST to subjective permissibility, my focus here is exclusively on objective permissibility. But I will clarify the obvious terms of art.

Morally better. I want COST to be acceptable to as broad an audience

²⁷ I concede that the foregoing argument might be consistent with a number of different principles besides the one I defend.

as possible, so it must be neutral on what makes one act morally better than another. It could have to do with the outcome that the act realises, or the act's intrinsic properties. One act might be better than another just in case it realises more agent-neutral value. Or we might rank acts by some complicated consequentialised version of common-sense morality, attending to both agent-relative and agent-neutral reasons.²⁸ Or there might be a more straightforward way to induce a deontological ranking on acts.²⁹ And (with a little adjustment) we could even view COST as being no more than a detailed specification of a particular *pro tanto* duty of beneficence, rather than as, for example, a criterion of right action.

Costs. Given the argument of section 2, clearly the costs in COST must concern the agent's self-regarding interests. This is controversial. Kagan [1989: 233ff], for example, argues that commonsense morality recognises options to favour one's own projects and commitments, even when these options cannot properly be understood as protecting one's self-regarding interests.³⁰ I think this is wrong. There are indeed agent-relative reasons to pursue particular relationships, commitments, and projects, but these do not (on their own) generate options. Instead, they have just the same moral standing as agent-neutral moral reasons: they contribute to making one act morally better than another, rather than in determining the costs to the agent. Options derive from the agent's self-regarding interests rather than her agent-relative reasons more generally, because she has special authority (of the relevant kind) only over the former kinds of interests. I do not have a special authority over the value realised by my pursuit of some commitment, except insofar as my pursuing it involves sacrificing my self-regarding interests for its sake. Of course, often the pursuit of such projects will require considerable sacrifice. So I will have options to pursue them or not on those grounds. But this is because of my authority over my self-regarding interests, not because the agent-relative reason in question is itself noninsistent.³¹

This makes for an interesting distinction in the kinds of reasons we might have for acting partially towards those we care about. Suppose, for example, that a trolley is headed towards my son, and I can save him only by diverting it down a sidetrack, where it will kill one other person. I think that diverting the trolley is the morally best thing to do, given my agent-relative reasons to protect my son (given some background assumptions). But now suppose that there are ten people on the sidetrack. Then I think that the morally best thing to do might be to let the trolley hit my son,

²⁸ E.g. Portmore [2011: 35].

²⁹ Hurley [2013]

³⁰ Kagan [1989: 241].

³¹ I mention one further point only to set it aside. COST implies that one should assess moral betterness and reasonable costs on an act-by-act basis. This might be wrong. Perhaps we should instead consider sequences of acts, or even whole lives (Cullity [2006]). That's an important debate, relevant to the objection in 5.B below, but not one I'm going to get into here.

even taking agent-relative reasons into account. Nonetheless, I don't think it would be reasonable to expect me to do that—the cost to me of letting my child die would be far too great. So I would still have the option to turn the trolley. But if the number on the sidetrack is great enough (and I don't know how much that would be), then I might be *required* not to turn the trolley. In this case, turning the trolley is morally so much better than protecting my son, that the cost to me is reasonable.

One heuristic to test whether partiality is justified by one's agent-relative reasons, or by the appeal to cost, is to ask whether we would praise a person who let his son die in this case. Would we regard this as a heroic form of self-sacrifice? Or as something that was morally wrong?

Marginal. We can distinguish between the absolute costs and benefits of an action, and its marginal (or opportunity) costs and benefits. The absolute costs can be measured by comparing the agent's self-regarding interests before and after the action; the absolute (moral) benefits are determined by comparing the world before and after the agent acted. The marginal costs and benefits are determined by comparing the act under consideration with the alternatives available to the agent.

Our authority over our own interests runs out when we are in a position to help others (or refrain from harming them) at no cost to ourselves. Inflicting unnecessary suffering is obviously wrong, and if benefiting others is costless, then it is morally required.³² This principle flows as naturally from individuals' moral status as does the justification of moral options. Though the reasons they can give us are not exhausted by the contribution of their interests to how the world goes, other people's interests obviously do matter, and if we can advance them at no cost to ourselves, we are required to do so.

And clearly marginal costs are what matters for this principle, not relative ones.³³ In the now-familiar case, the risk of entering the burning building might be such that I am permitted to remain outside if I want; but if I enter, I must save the child stranded inside, rather than my iPhone; indeed, if I can save two children instead of one, at no additional cost to myself, then it is wrong to save only one.³⁴ A similar phenomenon arises in the ethics of self-defence.³⁵ If I can save myself from a threat posed by an attacker either by killing her or by wounding her in the leg, then no

³² I think that Quinn [1989a]; Woollard [2013] would have to deny this, which is one problem with focusing on our authority over our bodies rather than our self-regarding interests.

³³ Curtis [1981: 311] is the only philosopher working on this topic to take a non-marginal view: 'if the cost or risk is considerably less significant [than the moral value of the end], the action is morally required; if considerably more significant, the action is foolish or unwise. But if the cost or risk is roughly *as* significant as the moral value of the end, the agent has done something which is "above and beyond the call of duty"—something which is morally good, but not morally required'.

³⁴ See, e.g. Kagan [1989: REF]; Portmore [2011: 140, 46-8].

³⁵ Lazar [2012].

matter how culpable she is, no matter how wrongly she has acted, I am prohibited from killing her—even though it would clearly be permissible if it was the only way I could save myself.

Or consider a further case: suppose that Alice could save 100,000 lives by sacrificing her own. If her only options are inaction or self-sacrifice, then let's stipulate that she's required to sacrifice her life—the number of lives saved is great enough that the cost to her, though maximal, is not unreasonable. But now suppose a third option becomes available: she can save 99,999 lives at the cost of a scratch to her finger. Even if 100,000 lives matter enough to make her clearly required to sacrifice her own life, if there are no other alternatives, if she has another option that saves only one less person, at a much lesser cost to herself, then she cannot be required to die. Saving 99,999 is morally worse than 100,000 (a life is a life!) but it is unreasonable for Alice to bear that additional marginal cost (her own death, vs. a scratch to her finger) for the sake of one additional life.

Reasonable. We have a special authority over our own interests. In virtue of this authority, I am not required to sacrifice my own interests to degree X, just in case doing so will realise X+I for someone else. But this authority is not boundless. I *am* required to sacrifice my self-regarding interests when doing so realises a significant enough moral improvement. The first clause of COST permits me to prefer my own interests as long as realising the morally better alternative would impose an 'unreasonable' marginal cost on me.³⁶ For any given morally better alternative to the act under consideration, we need to know whether the additional moral benefits are great enough to require the agent to bear the additional costs it involves. This means establishing a kind of proportionality between the moral good and the costs to the agent.

I doubt whether we can say more about this than we can about other kinds of moral weighing. But to work out what costs the agent can reasonably be expected to bear, we must do more than simply ask how great the marginal cost is, and how significant the marginal moral benefit. We must also ask if there are reasons to limit the agent's authority over her own interests in this case. For example, if the agent is responsible for a situation arising in which her act will have adverse consequences for someone else, that responsibility might make it reasonable for her to bear a much greater degree of cost than if she were not responsible in that way.³⁷ Similarly, if her act will affect people with whom she shares special relationships, and to whom she has associative duties—her children, for example—then her authority over her own interests would again be restricted.³⁸

Clause (a) provides for both agent-favouring and some agent-sacrificing options. It allows us to act suboptimally when all the better alternatives have unreasonable costs. But it also permits us to sacrifice our

³⁶ The central idea here is the same as in Scheffler's [Scheffler [1994] 'hybrid view'].

³⁷ See Miller [2001].

³⁸ Hurka and Shubert [2012].

interests for the sake of the greater moral good. Supererogatory acts will not be morally outranked by alternatives that have reasonable expected costs. Clause (b) allows for a different kind of self-sacrifice—when the sacrifice actually makes the overall outcome morally worse, but only in virtue of costs that befall the agent. It has two key elements.

Every such reasonable alternative. For a self-sacrificing act to be permissible, there must be no alternative act for which each of the following is true: it has reasonable marginal costs to the agent; it is morally better; and its moral betterness does not reduce to a benefit to the agent. If there is even one act for which these three conditions are true, then the act of self-sacrifice is impermissible.

Costs borne by the agent. The agent has authority over her self-regarding interests, so if the alternative acts are better only in virtue of their impact on her self-regarding interests, then she is entitled to forego them. This is consistent with mere self-sacrifice, in which the difference between the two options is *only* the difference in cost to the agent. And it's consistent with altruistic self-sacrifice, in which the alternative chosen is better for someone else and worse for the agent, but the cost to the agent is greater than the benefit to the other person/s.

Of course, many think that we have self-regarding duties, which prohibit us from making some kinds of self-sacrifice.³⁹ Although I disagree, my position on this is quite radical, and COST should not depend on it. We could capture this in different ways; I think the most plausible is to argue that in some cases of self-sacrifice, the sacrificing option will be worse than the alternatives not merely in virtue of the costs to the agent, but because of her breach of her self-regarding duty, which would be an additional agent-relative or impersonal reason for her not to bring that outcome about, which would not be wholly reducible to the damage to her self-regarding interests.

My view, though, is more radical than this. I think there are no self-regarding duties. Indeed, I am tempted to go even further: even if my self-sacrifice has costs for others, as well as for me, it might still be permissible as long as the costs to them derive from my no longer providing them with a benefit to which they have no right. Suppose, for example, that Alice is a much-loved actor, whose death will make millions of people very sad. We might reasonably think that she is still morally permitted to take her own life, because those millions have no right that Alice stay alive just to avoid making them sad. However, perhaps cases like these can be catered for by clause (a). Alice's death is *not* outranked by another option that has reasonable marginal costs, because she cannot reasonably be expected to forego exercising her free choice in such a fundamental matter as whether she lives or dies just to ensure that her fans are kept happy.

³⁹ For some fervent advocates of self-regarding duties, Hampton [1993]; Hurka and Shubert [2012].

It should I hope be clear how COST derives from the argument of section 2. Clause (a) provides for agent-favouring options, as well as the option to sacrifice oneself for the greater good. Clause (b) provides for suboptimal self-sacrifice. Both clauses are driven by the agent's authority over her own interests: they both provide space for that authority, and limit it. The space is provided by allowing the agent not to maximise in these ways. The limits are provided by insisting that if some morally better alternative comes at a reasonable marginal cost to the agent, then she is required to take it. This limits both agent-favouring options, and suboptimal self-sacrifice, by allowing for the possibility of self-regarding duties.

Clause (a) also entails that, in any decision problem, there is at least one option that constitutes 'the least you can do'. This is the morally best act that has reasonable marginal costs relative to all morally worse alternatives. Anything better than that baseline involves going 'beyond the call of duty', as long as it does not involve gratuitous suboptimality. Anything beneath it is wrong—unless its costs fall exclusively on the agent, in which case it is licensed by clause (b).

COST puts the abstract arguments of section 2 into practice. It operationalises the idea that we are not mere sites for the realisation of value, but are instead independent beings, with authority over our self-regarding interests. But it is not the only way to operationalise those ideas. Does it fare better than the alternatives?

4. THE ALTERNATIVES

Most of the discussion of moral options has fallen into one of two brackets: some argue that options are indeed a basic element in commonsense morality; others argue that they are no objection to consequentialism. These debates are connected: the first group typically presses the objection that the second group answers. In this section, I discuss two broad techniques used by consequentialists to accommodate moral options.⁴⁰ The first retains everything about consequentialism except its decision rule: it rejects maximising, arguing instead that it is morally permissible to 'satisfice'. That is, an act is permissible if and only if it realises *enough* val-

⁴⁰ I leave aside a third, variants of 'hybrid utilitarianism', as developed by Scheffler [1994], Vessel [2010] and Sider [1993]. These involve ranking acts according to their overall utility, and then according to their utility for the agent, and saying that an act is permissible if no other act surpasses it on both rankings. To develop these accounts in such a way that they fit with the intuitive data, as for example Vessel [2010] tries to do, I think you need to make modifications that leave you with a more complicated version of COST, with one crucial difference: COST naturally makes space for the idea that various moral reasons might affect what counts as a reasonable cost for you to bear (for example, the fact that you yourself are responsible for that cost coming about). Versions of hybrid utilitarianism developed thus far have no room for this possibility, and would need to be further complicated to accommodate it.

ue (as we'll see, this is only a first pass). The second approach retains a maximising approach to *moral* reasons, but argues that sometimes it is *rationally* permissible to act in ways that are morally suboptimal. In the following subsections, I summarise these views, then indicate why mine improves on them. I'll assess them against three standards: fit, simplicity, and strength. The first covers how well each principle maps onto the intuitive data. The second asks whether, for example, it unduly multiplies categories and criteria. And the third refers to each principle's explanatory strength—does it connect up with powerful ideas that make sense of moral options? Although I think my approach fits the intuitive data somewhat better than the alternatives, I think that its main advantage is its explanatory power.

A. Satisficing

In recent years enough has been written to discredit a simple satisficing consequentialism that adding to the objections would simply be piling on.⁴¹ But Jason Rogers has recently proposed a modified version of the view, which, he argues, overcomes the standard objections from Bradley, Mulgan and others, that if morality permitted satisficing, then it would license not only the gratuitous failure to provide others with benefits (this is basically what defines satisficing), but also the gratuitous imposition of harms.⁴² For as long as an act is permissible just in case the situation it realises has more value than some threshold, circumstances can arise in which we actively destroy value, without lowering the situation below the relevant threshold.⁴³ Rogers proposes the following principle in response:

SAT: There is a number, n , such that: An act, A , performed by agent S , is morally right iff either (i) the value of the situation after A is at least n , and is at least as high as the value of the situation prior to A , and any overall better alternative to A , A^* , is such that: [were A^* to be enacted instead of A , either S 's resultant personal welfare level after the enactment of A^* would be marginally significantly less than it was prior to the enactment of A^* , or the value of the situation after the enactment of A^* would not be appreciably greater than the value of the situation after the enactment of A]; or (ii) A maximizes utility. Rogers [2010: 216]

It's worth splitting out the different elements. The first part of (i) pre-

⁴¹ For example, in Pettit [1984]; Mulgan [2005]; Bradley [2006].

⁴² Mulgan [2005]; Bradley [2006].

⁴³ Interestingly, this objection is structurally parallel to Kagan's similar complaint against other arguments for moral options—that they threaten to overgenerate, resulting in options to inflict harm.

cludes the agent from making the situation worse, even if doing so keeps the overall value above the threshold. This deals with the objection that satisficing licenses gratuitous harms. Then the phrase leading to the square brackets caters for the worry that satisficing permits gratuitous sub-optimality. The agent is required to maximise if doing so doesn't have severe costs relative to the good achieved, and if the good achieved is 'appreciable'. And then (ii) allows that it's always permissible to maximise utility.

Rogers thinks these modifications to standard satisficing consequentialism accommodate all the relevant intuitive data. I think this is wrong. SAT inadequately caters for agent-sacrificing options. If sacrificing my own interests lowered the value of a situation below n , then it would be impermissible. It's hard to tell, in the abstract, just how implausible this is, but it seems very likely to conflict with standard judgements of when agent-sacrificing options are permissible.

Even above the threshold, Rogers cannot cater for permissible self-sacrifice: suppose that I can choose whether to go (on my own) on holiday to Fiji or to Belgium. Going to Fiji will realise appreciably more utility for me than going to Belgium, though going to Belgium would be good enough. It doesn't involve a cost to my personal welfare—indeed, it makes me better off. So I'm required to go to Fiji. That seems a mistake.

SAT also fails to accommodate altruistic self-sacrifice. Suppose that I can choose either to go on holiday to Fiji myself, or to pay for your holiday to Belgium. Either option realises enough utility. But if I go to Fiji, then I'll enjoy it appreciably more than you will enjoy going to Belgium. Now suppose that I decide, altruistically, to sacrifice my trip to Fiji so that you can go to Belgium. Is that permissible? Well, no: there is an alternative act, which does not involve marginally significant costs to my personal welfare, and which is appreciably better. So I am again required to maximise. I do something morally wrong by sending you to Belgium instead of taking my own trip to Fiji. Again, this is the wrong judgement.

Besides failing to account for agent-sacrificing options, Rogers' principle also requires maximisation when it is impossible to reach the morally acceptable threshold. Maybe such situations cannot arise, but since Rogers [2010: 201] says that the threshold is absolute rather than relative, it does seem *prima facie* possible. And in these cases, only (ii) will apply. Which means the agent will be required to maximise utility. Which in turn means that she will have to sacrifice her own life, say, if by doing so she can save the life of one other person who will be only one util better off than her. I think this is no more plausible in these tragic situations, when nothing one does is good enough, than it is when there are satisfactory actions available.

As to the explanatory strength of SAT: the problem here, and we will see this again in section 4.B, is simply that there is no attempt to argue for moral options. The goal is simply to identify a principle that can accom-

moderate them. Reflective equilibrium then provides the glue that holds it all together. This is clear in Rogers [2010: 199, 208].⁴⁴ I think we should hope for more than a demonstration that it is *possible* to accommodate moral options. I want to know whether we *should* do so. And Rogers offers little guidance on that score.

SAT misses some important intuitive judgements. It lacks explanatory strength. And it is needlessly complex: the threshold not only generates problems that the first phrase of (i) has to resolve, but it is also otiose. If we remove the threshold, we get something like this: An act is permissible if and only if either (i) it maximises expected utility or (ii) it falls non-trivially short of a better alternative, but the additional costs of that alternative cannot reasonably be required of the agent given the additional benefits. Together those two clauses would be extensionally equivalent with clause (a) of COST. This revised version allows us not to worry about permitting harmful actions that drop us below the threshold, because the threshold is gone. Gratuitous harms are ruled out on the same grounds as gratuitously failing to provide benefits. A threshold could in theory help accommodate some self-sacrificing options, but, as we have just seen, the other modifications in SAT block it from capturing standard such options. COST (b) covers those options, without inviting the same objections as does the threshold.

What's more, COST can give a better account than can SAT of why we might find the idea of a threshold tempting. COST allows us to work out a non-arbitrary threshold, which constitutes the least you can do in a given situation: the morally best option that, relative to all worse options, has reasonable marginal costs. Recognising this threshold does not bring with it the problems of the satisficing threshold—it still rules out gratuitous suboptimality above the threshold, and it licenses various forms of self-sacrifice. We can capture all the virtues of SAT, without inviting its problems, in a simpler principle.

Ultimately I think that the move to satisficing is a technical fix, introduced to cater for some intuitive judgments that are hard to square with a maximising consequentialism. But the threshold is unmotivated, and causes more problems than it solves. We do better to ask first what grounds moral options, and then devise a principle that reflects that grounding. COST is that principle.

B. Morality and Rationality

The second approach is more popular than the first. As a result, there are more variants, and I cannot do justice to them all. They share the idea that we should distinguish what we *morally* ought to do from what we *rational-ly* ought to do, all things considered. On this view, the 'ought' of all things considered rationality attends not only to moral reasons, but also to pru-

⁴⁴ See also Slote [1984b: 157].

dential ones, and perhaps to others such as reasons of friendship, aesthetic reasons, reasons of neighbourliness and so on. We can account for moral options by arguing that it is sometimes rationally permissible to bring about a morally worse outcome.⁴⁵ The paradigm case, of course, would be one in which the morally best option goes severely against the agent's self-interest. Some go further than this, and say that it is sometimes rationally permissible to breach a moral requirement (Dorsey [2012: 11]; Gert [2014: 218]). Others (Bratman [1994]; Portmore [2003]) once argued that an act is permissible if and only if it is either morally or rationally best, but nobody now holds this view.⁴⁶

My favourite version of this approach says that an act is morally required if and only if it is both morally best and best all things considered, and an act is all things considered permissible if and only if no alternative act is *both* morally better and all things considered better (Portmore [2011: 4]). On this account, agent-favouring options are instances in which an act is morally outranked, but not all things considered outranked by some alternative. It would be morally better to do the alternative, but the cost to the agent successfully blocks this from becoming an all things considered requirement. Agent-sacrificing options, by contrast, are instances in which the agent-sacrificing option is all things considered outranked, but not morally outranked, by some alternative.

Any approach to options with this form has to distinguish very clearly between moral and non-moral reasons. This is no easy task. Portmore, for example, argues that a moral reason is a fact that, morally speaking, counts for or against some action.⁴⁷ He does not say what he means by 'morally speaking'. But he does identify at least one class of reasons that are *not* moral: the agent's self-interest is not, *for that agent*, a moral reason (it is, of course, a moral reason for others).⁴⁸ And this move is crucial. Without it, the appeal to rationality would not preserve important judg-

⁴⁵ Curtis [1981]; Wolf [1982]; Slote [1991]; Kagan [1994: 337]; Vessel [2010]; Portmore [2011]; Dorsey [2012].

⁴⁶ Hurka and Shubert [2012: 10] are pushing at an open door when they criticise this view.

⁴⁷ 'Moral reasons either have some moral requiring strength or, if they do not, they are mere moral enticers. Moral enticers can make doing what they entice us to do supererogatory, but they cannot make doing what they entice us to do obligatory. Thus, a moral reason is a reason that, if sufficiently weighty, could make an act either obligatory or supererogatory. A reason that could only justify—that is, a reason that could not make an act obligatory or supererogatory but could only make an act permissible—would be a (morally relevant) non-moral reason.' Portmore [2011: 123].

⁴⁸ '[T]here is nothing, morally speaking, that counts in favor of promoting one's self-interest per se. This is not to say that one never has a moral reason to do what will further one's self-interest; one often does, as when doing one's moral duty coincides with promoting one's self-interest. The claim is only that the mere fact that performing some act would further one's self-interest does not *itself* constitute a moral reason to perform that act, for the mere fact that performing some act would be in one's self-interest is never by itself sufficient to make an act obligatory, or even supererogatory.' Portmore [2011: 128]. See also Portmore [2011: 96, fn 39]

ments about agent-sacrificing options.⁴⁹ If we had moral reasons to pursue our self-interest, then the appeal to rationality would not preserve agent-sacrificing options, because the non-sacrificial alternatives would be both morally and all things considered better.

The success of the appeal to rationality, then, depends on its insistence that the agent's self-interest cannot make an outcome morally better. And this is its greatest weakness. Although I am not *merely* a site for the realisation of value, I am *also* such a site! My interests are not the only thing about me that matters morally—but they do matter morally, even if I am the one acting. The contrary view deprives us of a whole species of justification: according to Portmore, one act simply cannot morally outrank another *in virtue of* contributing to the agent's self-interest. So if I choose an option that favours my interests, it is not (at least not for that reason) morally better than the self-sacrificing alternatives. It might still be rationally permissible. But it is morally worse.

I think this is a mistake. Sometimes I have a moral justification for acting in my own interests, not merely a rational one. What's more, this is not just a question of terminology: Portmore's approach undermines interpersonal justification. For me to justify taking the option that benefits me, I have to say that though morally worse, it's the all things considered rational choice for me to take. If I were the party adversely affected by this decision, I'd be inclined to say 'well **** you and your all things considered rationality!'

This approach also deals poorly with some important cases. Suppose we find some manna, which will bring me 100 units of happiness and you 10 units. If I claim it, Portmore's view would say that I am bringing about a worse outcome, though one that is all things considered justified by the force of my prudential reasons to gain the 100. In other words, I am acting on a kind of right to be selfish. This seems wrong: I am realising a much better outcome, and that should matter.

Or suppose that it's not manna, instead it's resources that I have earned through my own labour and ingenuity. If I claim those resources rather than giving them up to you, I am not simply acting out of permissible self-interest, I am taking what is mine by right (given some background assumptions).

Or now suppose that I can choose between two distributions, A and B. A has everyone equal, while B gives everyone else a little more, and me much less. Again, on Portmore's view I am permitted to bring A about because B is not all things considered rational, but still, B is the morally better distribution. I think that is wrong: the morally better distribution is the egalitarian one, even though it is better than the other only in virtue of how my self-regarding interests are satisfied.

Of course, Portmore has a response handy: he can argue that there are

⁴⁹ Sider [1993]; Vessel [2010].

independent moral reasons, not having to do with my self-regarding interests, to ensure that manna is used efficiently, or that people get the results of their labour, or that distributions are equal. But this response has its own problems: if those reasons make outcomes morally better, then I am morally *required*, in each case, to insist on my rights. I am required to take the manna, the resources, and secure the egalitarian distribution. But this is wrong; in each case I am entitled to forego these benefits to myself if I want to. Even if a principle of justice confers on me a right to some good, I have a right to give it up if I want to.

Not counting self-regarding interests as moral reasons also has troubling implications for the ethics of self-defence. Suppose a culpable attacker threatens the life of an innocent defender. The standard justification for killing in self-defence goes like this: someone has to die, and the defender must choose between killing and being killed.⁵⁰ There is a strong presumption against killing another person to save oneself, grounded in the distinction between doing and allowing. To override that presumption, we need a moral asymmetry between the two people whose lives are at stake. In this case, if the defender saves herself, she saves an innocent person; if she kills the attacker, she kills someone who is not only culpable, but culpable for this very situation arising. The fact that the defender can save her own life is therefore crucial to justifying killing the attacker. But if the defender's survival cannot make the outcome morally better, because she does not have moral reasons to act in her own interests, then what can justify lethal defensive force?

Again, Portmore might respond that even if killing the attacker is not morally justified, it is all things considered rational. But even if this gets the right deontic verdict, it does so for the wrong reasons. Killing a culpable attacker in self-defence is *not* a morally bad outcome that one is entitled to bring about because the cost to you of not doing so is just too great. The defender acts justly, and brings about a better outcome, when she saves her own life. The culpable attacker's interests are discounted by his culpability; the defender is innocent, so her interests are not discounted.

Again, one could argue that the defender has reasons of justice to kill the culpable attacker, but that would again imply that she is required to do so, which is clearly false—if she wants to let herself be killed, she may. What's more, it also implies that there is some positive reason for her to kill the attacker—like retribution for his wrongdoing. And most people think that retribution (and indeed desert, to which it responds) have no place in a plausible theory of self-defence. The point is not that 'it's a good thing' to kill the culpable attacker. It's simply that *it is better* to kill him than let herself be killed.⁵¹

⁵⁰ For a canonical view, see e.g. McMahan [2005].

⁵¹ The same problem applies when considering the infliction of proportionate harm on the innocent as a side-effect of saving oneself from some dangerous threat. Suppose that the defender can save herself only by tossing a grenade, which is likely to injure but not

Strictly speaking, Portmore's sole aim in developing his principle is to get the right deontic verdicts. But it matters not only what those verdicts are, but how they are reached. And arguing that ϕ ing is permissible because it is morally best is quite different from arguing that, though it is not morally justified, it is all things considered rational. The agent's self-interest gives moral reasons in its own right, and plays a crucial role in the application of a number of other important principles.

The appeal to rationality cannot adequately account for agent-sacrificing options. That is its principal failing. But it has other flaws too. First, I am sceptical about whether our intuitions about morality versus all things considered rationality are reliable. The key problem is fineness of grain. While I think we can plausibly appeal to intuition to deliver verdicts on all things considered permissibility, I cannot see how intuition can deliver verdicts on whether an action is permissible *because* it is morally best, or because it is rationally best. The appeal to rationality relies on our being able to pull apart multiple different senses of permission, appropriate to each of the different normative spheres that it invokes. I think any such intuitions are too theory-laden to do much probative work.

What's more, even this move's advocates think that some of its implications are counterintuitive. On their view, for example, performing supererogatory acts is irrational. Dorsey [2013: 372] is relatively unperturbed, but insofar as I can make out my pretheoretical views on this, I agree with Kagan and others that this is a genuine cost.⁵² And it is one that we do not have to bear. According to COST, supererogatory acts involve costs that exceed what can reasonably be expected. It might be rational to bear those costs, but they are severe enough that it is the agent's decision whether to do so. Portmore [2011: 54] argues that in these cases one's reasons of self-interest 'successfully counter' the moral reasons. They outweigh them. My view is different: I think those costs give you a right to refuse. I think this better tracks the intuitive data, such as they are, about rationality.

So much for 'fit'. On to explanatory strength. And again I think the appeal to rationality again looks like a technical fix, introduced to preserve sensible deontic verdicts, but without intrinsic motivation. Portmore is aware of this charge and resists it, but he does so by tying his very abstract principle to principles that are *even more* abstract. In the end each adherent of this position admits that their basic approach is that of reflective equilibrium—they want to find a simple principle that maps onto the

kill an innocent person standing near the culpable attacker. If the defender is justified in doing so, it must be because the outcome in which her life is saved is morally better than the one in which she is killed, and the bystander avoids that additional harm. To justify her action in the right way, we have to include the defender's interests in the calculation of proportionality.

⁵² Kagan [1994: 349]; Hurka and Shubert [2012].

intuitive data.⁵³ That is fine, as far as it goes. But my approach ties moral options to some basic features of our morality—to our moral status, and the authority we have over our own interests. Instead of relying on a definitional claim—that reasons of self-interest cannot be moral reasons—it gives an explanation for *why* those reasons cannot generate moral requirements, because we have a fundamental authority over our self-regarding interests.

My approach is more concrete, and as a result more readily operationalisable than the appeal to rationality. It might not be simpler, but then if you were to apply, say, Portmore’s principles, I think that in practice you would operationalise them by thinking about the questions that COST requires you to ask. To work out whether some option is all things considered required, you would need to work out whether it involves exacting a reasonable cost for the agent. The appeal to rationality gives us a highly abstract family of principles, which are functionally coextensive with COST, while COST is both more concrete, and more informative.

So, although the appeal to rationality can accommodate agent-favouring options, its means of accounting for agent-sacrificing options is unconvincing, and it is again too free-floating an attempt to cobble together a version of consequentialism that can map onto our considered judgements. We need to know why costs to the agent can block moral requirements, not only what follows if they can. And we need to know why the agent’s self-regarding interests cannot generate moral requirements, not only to define our terms in such a way that they do not.

5. OBJECTIONS

Thus far I have given an argument for moral options, and shown how it meets the basic desiderata of justifying both kinds of options, as well as neither overgenerating options, nor generating requirements instead of options. I have formulated a principle that operationalises these abstract ideas, making them concrete. And I have shown how my principle and argument fit the intuitive data better than the two main alternatives, as well as being simpler, and having more explanatory power. But it obviously would not be fair to stop here. The argument from authority must face its own share of scrutiny.

A. Dorsey’s Objection

I will dwell most on an objection from Dorsey [2013: 367], which is in effect a more challenging iteration of Kagan’s concern that arguments for options will tend to overgenerate Kagan [1989: 22]. The objection starts with

⁵³ Gert [2000: 238]; Vessel [2010]; Portmore [2011: 11, 119 fn 1, 50]; Dorsey [2012: 12, 2013: 356]. Scheffler [1994] is an obvious exception.

a case.

Gus can choose either to spend his savings on a nice new car, or to donate them to a charity, who will use them to save 10 lives (with no further consequences). Suppose that he has the agent-favouring option to buy the car, rather than save the lives.⁵⁴ But now a different situation arises—let's not worry about the details—in which Gus can save 10 lives if he beats up Jerry. Even if doing harm is worse than allowing it, surely the stakes are high enough that the presumption against the former is overridden. So it is permissible to beat up Jerry, ensuring that the 10 are saved. But now, suppose that if Gus wants to buy his nice new car, he has to beat Jerry up first; if he refuses to beat Jerry up, he loses the savings that would get him the nice car. Surely we would all think that he *doesn't* have the agent-favouring option to batter Jerry so that he gets to buy the car.

Call 'beat up Jerry' J, 'buy the car' B, and 'save the 10' S. And use $X \geq Y$ to mean 'is permitted to choose X rather than Y', or 'X is weakly morally preferred to Y'. In the choice between B and S, $B \geq S$. In the choice between S and J, $S \geq J$. So since $B \geq S$, and $S \geq J$, we might think that, by transitivity, $B \geq J$. If 'buy the car' is weakly morally preferred to 'save the 10', and 'save the 10' is weakly preferred to 'beat up Jerry', shouldn't 'buy the car' be weakly preferred to 'beat up Jerry'? This would obviously be counterintuitive.

This objection rests on one assumption: that there is a single dimension of normative strength, so that if the reasons for A, B, and C are such that you weakly prefer A to B, and B to C, then by transitivity you must also weakly prefer A to C. Dorsey makes this explicit at a number of points, emphasising that 'it is at least as morally important (that is, as important from the perspective of first-order moral reasons) to save ten as it is not to beat up Jerry' (Dorsey [2013: 367]), so if Gus has the option to buy the car rather than save the ten, he has the same option to beat up Jerry. Again: 'If agential sacrifice (by whatever mechanism) morally justifies pursuit of one's own interests against an action supported by moral reasons of strength s , it should also morally justify pursuit of one's interest against any other action supported by reasons of strength s - or-weaker.' (Dorsey [2013: 367]).

This assumption is false. Frances Kamm showed this in her Kamm [1985], where she distinguished between the 'efforts standard' and the 'precedence standard' in comparing moral reasons.⁵⁵ The 'efforts standard' governs what degree of cost a reason can require the agent to bear. The 'precedence standard' governs how that reason weighs against other moral reasons, when they compete. I will rename these, respectively, the

⁵⁴ If you find this implausible, change the number of lives saved to suit.

⁵⁵ There is an interesting analogy between these two standards and the 'justifying' and 'requiring' strength of normative reasons (Gert [2003]). In fact, I think that requiring strength is the same as stringency. Perhaps so there are really three dimensions of normative strength, not two: justifying, requiring/stringency, and seriousness.

stringency and the *seriousness* of a moral reason. Kamm gives a nice example, which is directly analogous to Dorsey's. Here is a version of it.

On my way to an appointment, I am alerted that a friend is urgently in need of a kidney transplant. I am permitted to break my appointment to donate my kidney to the friend. Now suppose that I have no appointment, but am planning to play golf on the afternoon when I hear about my friend's urgent need. I'm clearly not *required* to give my kidney to my friend, so I'm permitted to go to the golf club instead of to the hospital. And now, last, suppose that I am on my way to my appointment, when I get a hankering to play golf. Clearly I'm *not* permitted to break my promise, just to get in a quick nine holes.

If Dorsey was right, then this set of very plausible judgements would be unintelligible, because we have already established that the reason to play golf has enough normative strength to ground a permission not to donate the kidney, and the reason to donate your kidney has enough strength to ground a permission not to keep your appointment, so it would be 'utterly arbitrary' (Dorsey [2013: 369]) to deny that your reason to play golf has enough normative strength to ground a permission to break your appointment.

Return to Dorsey's case. Gus's reason to save 10 is not sufficiently stringent to require him to forego buying the new car. Saving 10, however, is sufficiently serious to override Gus's reason not to beat Jerry up. In the competition between his reasons to save the 10, and his reasons not to beat up Jerry, the former set of reasons take precedence. But then beating up Jerry is clearly sufficiently stringent to require Gus to forego his new car rather than assault Jerry. The three two-place comparisons in Dorsey's example do not all invoke the same dimension of normative strength. So there is no reason to expect transitivity.

Dorsey might respond that it's all very well to posit a distinction between different kinds of normative strength, and to show that it applies in other cases, but if it is just an unmotivated technical fix the objection still stands. Properly motivating the distinction would require a paper in its own right, but here is one way that paper could go.

The basic question is: why don't the stringency and seriousness of moral reasons co-vary? Why don't they both derive from a deeper idea of 'moral importance'? The answer has to do with whether those reasons involve a significant agent-relative component.

Agent-relative reasons are indexical; they unavoidably include some pronominal back-reference to the agent to whom they apply.⁵⁶ Often the same fact will generate both agent-relative reasons and non-indexical, agent-neutral reasons. If my son is drowning, then I have both agent-relative (he's my son!) and agent-neutral reasons (he's a human being!) to save him. The more prominent the agent-relative component of a reason,

⁵⁶ Pettit [1987]; McNaughton and Rawling [1991, 1995].

the more likely it is that stringency and seriousness will come apart. If the agent-relative dimensions of that reason are particularly strong, then it is likely to be particularly stringent. After all, this reason especially applies to *you*, so it can more readily justify imposing costs on you than it can justify imposing them on others. Of course, agent-relative reasons *can* justify imposing costs on other people besides the agent. But if the reason has particular agent-relative force, then that is more obviously grounds for the agent bearing additional costs in order to act upon it, than for his imposing those costs on others.

Gus's reason not to harm Jerry has a prominent agent-relative component; his reason to save the 10 does not. If Gus could harm Jerry to prevent two identical harms being inflicted by someone else, he ought not do so (this shows the agent-relative component of the reason). But if Gus foregoing saving the 10 would ensure that 20 were saved by someone else, then he should forego saving the 10. In Kamm's case, the reason to keep your appointment is grounded in a promise, so is agent-relative.⁵⁷ And we can generate very similar cases using our associative duties to those we care about, which are agent-relative too. So, the distinction between the seriousness and stringency of a moral reason is both well-motivated and fits the intuitive data. And it defuses Dorsey's objection.

In fact, Dorsey's objection poses more of a problem for his view than it does for his targets'. After all, since he does not see a distinction between stringency and seriousness, he must believe that it is all things considered rationally permissible for Gus to beat up Jerry rather than miss out on the new car; indeed, it would be irrational for him to sacrifice his new car in order to avoid making the assault. This is a bitter pill to swallow.

B. Did Gandhi Act Wrongly?

The next objection has been made before by others, but was pressed on me with particular force by Josef Holden.⁵⁸ Does my view really imply that Gandhi, for example, acted impermissibly if on some occasion he could have done *slightly* more good, at little or no additional cost to himself? Can we really say that people who do *so much* more good than anyone else, at so much cost to themselves, are acting impermissibly when they miss the chance to save an extra life, or do a little more good?

The right response to this objection is first to deny the underlying intuition, at least in a subset of cases; and then to show that in other cases a sensible interpretation of COST will deliver intuitively plausible results.

On the first point: suppose that Gandhi had an opportunity to realise a somewhat morally better outcome at absolutely no cost to himself. I think

⁵⁷ Example: I can save my child or save forty others. Clearly I'm permitted to save my child. I can save forty children through action that will collaterally kill four people. Probably I'm permitted to save the forty, killing the four. I can save my child only through action that will collaterally kill four people. I'm obviously not permitted to do so.

⁵⁸ See, for example, Splawn [2001].

missing that opportunity *would* be impermissible. Choosing a worse outcome when a better one is available, and is costless to you, is much the same as inflicting unnecessary harm. Obviously we need to be sure we have the right account of what makes one outcome better than another (which will take into account agent-relative reasons, for example). But once we have that in place, gratuitous sub-optimality is no better than gratuitous suffering. So yes, if Gandhi had passed on an opportunity to do more good at no cost to himself, that would have been wrong. I'm not saying that *I* would upbraid him for it—but that is a question of standing and hypocrisy, not one that speaks to the underlying moral phenomena.

However, the idea that underlies the objection does make sense, and can be put to some use. When someone has already given *so much* of themselves to promoting the good, it becomes unreasonable to expect them to do much more. My account should accommodate this datum. I need to argue, then, that what counts as a 'reasonable cost' depends in part on how much cost you have already taken on. As the agent takes on more and more cost, the additional increment in moral goodness needed to require him to take on an additional unit of cost increases. I won't offer a full defence of this view, but it seems plausible.

Interestingly, this line would offer an alternative to 'whole-life' approaches to supererogation and duties of beneficence, according to which we should calculate what cost you are required to bear in a given instance by looking at your opportunities to do good over your life as a whole.⁵⁹ My proposal is intermediate between that and an ahistorical focus on the act alone. It can evaluate actions independently of considering a person's life as a whole, which is a virtue. But it allows us to take history into account, so some sacrifices might be required of those who have thus far led selfish lives, but merely optional for those who have already done a great deal.

C. What About Essentially Non-Requiring Reasons?

Some might think that COST is too narrow to do justice to the phenomenon of supererogation. They might think that commonsense morality recognises cases in which helping others is optional even though it is costless to the agent. Horgan and Timmons [2010], for example, hold that at least some actions are supported exclusively by 'favouring' rather than 'requiring' reasons, such that one need not be required to act on them even if they involve no sacrifice at all. Typical examples include giving directions to someone who is lost (common enough around Coombs), or taking a stranger on a fun outing (ditto).⁶⁰

I think we can resolve these objections in two ways: first, by recognising that in any balance of moral reasons, it is hopeless to expect precision, so when very small costs and benefits are at stake, we should be sceptical

⁵⁹ See, for example, Cullity [2004]; Portmore [2011].

⁶⁰ See also Kagan [1989: 243]; Dorsey [2013: 375].

of sharp boundaries and clean verdicts.⁶¹ It might, then, be permissible not to provide trivial benefits simply because morality is imprecise, and the benefits are trivial enough that we are in the penumbra of permissibility.

But insofar as we stipulate imprecision or indeterminacy away, I think this objection appeals to an empty set of cases. Giving someone directions (especially in Coombs) takes time, as does taking someone on a fun outing. Even if you enjoy the outing, there are still other things that you could be doing, so some cost is involved. As long as the benefits to the beneficiary are trivial, these cases already fall under the remit of COST: the moral improvement is small enough that even the slight self-sacrifice involved is not one that can reasonably be required.

And if we then stipulate the cost away, and argue that there is *genuinely* no cost to the agent, then I simply deny the intuition. If I can benefit someone else at no cost to myself, then what could possibly justify my not doing so? Think again of harms. Suppose that I have the option to misdirect someone—send them off around the wrong hexagon of Coombs, for example. Even if the cost to them is trivial, if there is zero gain to me in doing so, what could possibly justify being such an ass?

D. Ratcheting Up?

Any view that emphasises the importance of comparing each act with the available alternatives will face the worry that we might change the deontic status of some act by introducing a number of alternatives, which together ‘ratchet up’ what you are required to do, resulting in an act that was antecedently merely permissible becoming required.⁶² Suppose, for example, that we use numbers to indicate [overall moral betterness, cost to the agent]. And suppose that the agent starts out with these options:

A [2, 0]
E [10, 4]

Let us stipulate that the marginal cost to the agent (4 units) is severe enough relative to the 8 units of moral gain that the agent cannot be required to choose E over A. But now suppose that the following options appear:

B [4, 1]
C [6, 2]
D [8, 3]

And let us stipulate that, in each case, the additional two units of moral betterness are weighty enough to ensure that the additional one unit of cost is required. In other words, A, B, C and D are all now impermissible,

⁶¹ Dougherty [2013].

⁶² Thanks to Kai Spiekermann and Robbie Williams here.

since E is morally better than them, and by stipulation involves bearing a reasonable cost relative to D, which is reasonable relative to C, and so on. If only A and E were available, both A and E would be merely permissible. But the presence of B-D, which are themselves impermissible options to take, renders E required and A prohibited. We might think this odd.

I think we should not worry too much about this kind of objection; according to my view (and most of those I have criticised), sometimes whether an option is permissible *will* depend on which other options are available. However, if you find this disconcerting, notice that we can generate the puzzle only by stipulating that though a 2:1 ratio of moral betterness to personal cost is enough to render the move from A to B, B to C, C to D, and D to E required, an 8:4 ratio is not enough to make the move from A to E required. Obviously it is silly to think we could put precise numbers on our moral reasons, but insofar as we can model weights using numbers as a representational device, and insofar as you find the worry raised by the case compelling, it merely indicates an interesting formal constraint on how moral betterness should weigh against cost to the agent: the ‘reasonableness ratio’ should be constant regardless of absolute stakes. If 2:1 is enough to make an alternative required, then 8:4, or 200:100, should do the same.

Does this conflict with my response to the Gandhi objection? I think not. The amount of cost one has *already borne* might alter the amount of cost one can now reasonably be required to bear, for the sake of a given increment of moral betterness. But this is consistent with the ‘reasonableness ratio’ being constant at any given time, in any given decision, which is all we need to avoid the counterintuitive implications of the present objection.

6. CONCLUSION

Moral options are an ineliminable element of commonsense morality. But not enough has been said to justify them, in the face of Kagan’s onslaught of criticisms in the late 1980s. Unfortunately, Kagan’s critique shaped the subsequent defences, which sought to resurrect arguments that he had eviscerated—whether those drawing on our unfortunate inability to motivate ourselves to maximise, or others emphasising the importance of preserving space for our projects and commitments. Even the practice of viewing moral options as an embarrassment for consequentialism owes much to Kagan’s framing of the debate. The result is a paucity of compelling arguments for moral options, and a widespread tendency to view this as a technical challenge for consequentialists. In this paper, I have tried to reverse those two trends: to make a compelling case that moral options are grounded in a fundamental principle of justice; and to offer a principle operationalising those options that is better motivated, simpler, and

more consistent with the intuitive data than the available alternatives. We are permitted to prefer our own interests even when sacrificing them would realise a greater good, and to sacrifice our interests even when preserving them would make the world better, because they are *our interests*. We have a special authority over them which, though not boundless, inscribes at the heart of morality our status as independent beings, each with our own lives to lead. Undoubtedly having these options is good for us, and perhaps we are incapable of acting as though we didn't have this authority. But that is not *why* we have these options. They are instead grounded in a basic principle of justice, and a fundamental idea of moral status. Because we have moral status, our interests matter. But because we have moral status, *we matter too*.

REFERENCES

- Bradley, B. 2006. 'Against Satisficing Consequentialism,' *Utilitas* 18/2: 97-108.
- Bratman, M. E. 1994. 'Kagan on 'the Appeal to Cost',' *Ethics* 104/2: 325-32.
- Broad, C. D. 1971. 'Self and Others,' in ed. David R Cheney, 1-9.
- Brock, D. W. 1991. 'Defending Moral Options,' *Philosophy and Phenomenological Research* 51/4: 909-13.
- Cullity, G. 2004. *The Moral Demands of Affluence*, Oxford: Oxford University Press.
- Cullity, G. 2006. 'Moral Demands of Affluence,' 1-297.
- Curtis, B. 1981. 'The Supererogatory, the Foolish and the Morally Required,' *The Journal of Value Inquiry* 15/4: 311-18.
- Dorsey, D. 2012. 'Weak Anti-Rationalism and the Demands of Morality,' *Noûs* 46/1: 1-23.
- Dorsey, D. 2013. 'The Supererogatory, and How to Accommodate It,' *Utilitas* 25/03: 355-82.
- Dougherty, T. 2013. 'Vague Value,' *Philosophy and Phenomenological Research* 89/2: 352-72.
- Gert, J. 2000. 'Practical Rationality, Morality, and Purely Justificatory Reasons,' *American Philosophical Quarterly* 37/3: 227-43.
- Gert, J. 2003. 'Requiring and Justifying: Two Dimensions of Normative Strength,' *Erkenntnis* 59/1: 5 - 36.
- Gert, J. 2007. 'Normative Strength and the Balance of Reasons,' *The Philosophical Review* 116/4: 533-62.
- Gert, J. 2014. 'Moral Rationalism and Commonsense Consequentialism,' *Philosophy and Phenomenological Research* 88/1: 217-24.
- Hampton, J. 1993. 'Selflessness and the Loss of Self,' *Social Philosophy and Policy* 10/01: 135-65.

- Heyd, D. 1982. *Supererogation*, Cambridge: Cambridge University Press.
- Horgan, T. and M. Timmons 2010. 'Untying a Knot from the inside Out: Reflections on the "Paradox" of Supererogation,' *Social Philosophy and Policy* 27/02: 29-63.
- Hurka, T. and E. Shubert 2012. 'Permissions to Do Less Than the Best: A Moving Band,' in *Oxford Studies in Normative Ethics, Volume 2*, ed. Mark Timmons, Oxford: Oxford University Press: 1-27.
- Hurley, P. 1995. 'Getting Our Options Clear: A Closer Look at Agent-Centered Options.'
- Hurley, P. 2013. 'Consequentializing and Deontologizing: Clogging the Consequentialist Vacuum,' in Oxford: Oxford University Press.
- Jackson, F. 1991. 'Decision-Theoretic Consequentialism and the Nearest and Dearest Objection,' *Ethics* 101/3: 461-82.
- Kagan, S. 1989. *The Limits of Morality*, Oxford: Clarendon Press.
- Kagan, S. 1994. 'Defending Options,' *Ethics* 104/2: 333-51.
- Kamm, F. M. 1985. 'Supererogation and Obligation,' *Journal of Philosophy* 82/3: 118-38.
- Kamm, F. M. 1992. 'Non-Consequentialism, the Person as an End-in-Itself, and the Significance of Status,' *Philosophy and Public Affairs* 21/4: 354-89.
- Lazar, S. 2012. 'Necessity in Self-Defense and War,' *Philosophy & Public Affairs* 40/1: 3-44.
- McMahan, J. 2005. 'The Basis of Moral Liability to Defensive Killing,' *Philosophical Issues* 15/1: 386-405.
- McNaughton, D. and P. Rawling 1991. 'Agent-Relativity and the Doing-Happening Distinction,' *Philosophical Studies* 63/2: 167-85.
- McNaughton, D. and P. Rawling 1995. 'Value and Agent-Relative Reasons,' *Utilitas* 7/1: 31-47.
- Miller, D. 2001. 'Distributing Responsibilities,' *Journal of Political Philosophy* 9/4: 453-71.
- Mulgan, T. 2005. *The Demands of Consequentialism*, Oxford University Press.
- Pettit, P. 1984. 'Satisficing Consequentialism,' *Proceedings of the Aristotelian Society* 58: 165-76.
- Pettit, P. 1987. 'Universalizability without Utilitarianism,' *Mind* 96/381: 74-82.
- Pettit, P. 1997. 'The Consequentialist Perspective,' in *Three Methods of Ethics*, ed. Philip Pettit, Marcia Baron, and Michael A. Slote, Oxford: Blackwell: 92-174.
- Portmore, D. W. 2003. 'Position-Relative Consequentialism, Agent-Centered Options, and Supererogation,' *Ethics* 113/2: 303-32.
- Portmore, D. W. 2011. *Commonsense Consequentialism: Wherein Morality Meets Rationality*, Oxford: Oxford University Press.
- Quinn, W. S. 1989a. 'Actions, Intentions, and Consequences: The Doctrine of Doing and Allowing,' *Philosophical Review* 89: 287-312.

- Quinn, W. S. 1989b. 'Actions, Intentions, and Consequences: The Doctrine of Double Effect,' *Philosophy and Public Affairs* 18/4: 334-51.
- Railton, P. 1984. 'Alienation, Consequentialism, and the Demands of Morality,' *Philosophy and Public Affairs* 13/2: 134-71.
- Rogers, J. 2010. 'In Defense of a Version of Satisficing Consequentialism,' *Utilitas* 22/02: 198-221.
- Scheffler, S. 1992. *Human Morality*, Oxford: Oxford University Press.
- Scheffler, S. 1994. *The Rejection of Consequentialism*, Oxford: Clarendon Press.
- Sider, T. 1993. 'Asymmetry and Self-Sacrifice,' *Philosophical Studies* 70/2: 117-32.
- Sidgwick, H. 1981. *The Methods of Ethics*, Indianapolis: Hackett.
- Slote, M. 1984a. 'Morality and Self-Other Asymmetry,' *The Journal of Philosophy* 81/4: 179.
- Slote, M. 1984b. 'Satisficing Consequentialism,' *Proceedings of the Aristotelian Society* 58: 139-63.
- Slote, M. 1991. 'Review: Shelly Kagan's the Limits of Morality,' *Philosophy and Phenomenological Research* 51/4: 915-17.
- Smart, J. J. C. and B. Williams 1973. *Utilitarianism; for and Against*, Cambridge: Cambridge University Press.
- Sobel, D. 2007. 'The Impotence of the Demandingness Objection,' *Philosophers' Imprint* 7/8: 1-17.
- Splawn, C. 2001. 'The Self-Other Asymmetry and Act-Utilitarianism,' *Utilitas* 13/03: 323-33.
- Stocker, M. 1976. 'Agent and Other: Against Ethical Universalism,' *Australasian Journal of Philosophy*.
- Suikkanen, J. 2014. 'Consequentialist Options,' *Utilitas* 26/3: 276-302.
- Tadros, V. 2011. *The Ends of Harm: The Moral Foundations of Criminal Law*, Oxford: Oxford University Press.
- Thomson, J. J. 2008. 'Turning the Trolley,' *Philosophy and Public Affairs* 36/4: 359-74.
- Urmson, J. O. 1958. *Saints and Heroes (W:) Ai MeIden (Ed.) Essays in Moral Philosophy*,
- Vessel, J.-P. 2010. 'Supererogation for Utilitarianism,' *American Philosophical Quarterly* 47/4: 299-319.
- Waldron, J. 1994. 'Kagan on Requirements: Mill on Sanctions,' *Ethics* 104/2: 310-24.
- Wolf, S. 1982. 'Moral Saints,' *Journal of Philosophy* 79/8: 419-39.
- Woollard, F. 2013. 'If This Is My Body...: A Defence of the Doctrine of Doing and Allowing,' *Pacific Philosophical Quarterly*.