

Moral Status and Agent-Centred Options

SETH LAZAR
Australian National University

Abstract

If we were required to sacrifice our own interests whenever doing so was best overall, or prohibited from doing so unless it was optimal, then we would be mere sites for the realisation of value. Our interests, not ourselves, would wholly determine what we ought to do. We are not mere sites for the realisation of value—instead we, ourselves, matter unconditionally. So we have options to act suboptimally. These options have limits, grounded in the very same considerations. Though not *merely* such sites, you and I are *also* sites for the realisation of value, and our interests (and ourselves) must therefore sometimes determine what others ought to do, in particular requiring them to bear reasonable costs for our sake. Likewise, just as my moral status grounds a requirement that others show me appropriate respect, so must I do to myself.

I. INTRODUCTION

Morally speaking, perhaps I shouldn't be writing this essay. I could be doing something much better with my time. I'm not skilled, but I'm a reasonably fast learner—perhaps with dedication and retraining 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 strange 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.¹ 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 Star Wars figures (for my son... mostly). But I'm not going *that* far wrong.

To give it a technical name, I think that I have a *self-favouring option to act sub-optimally*,

¹ Some advocates of indirect or subjective consequentialism argue that actually just living our lives in a more or less ordinary way *does* maximise (expected) value. See, for example, P. Pettit, 'The Consequentialist Perspective', *Three Methods of Ethics*, ed. Philip Pettit, Marcia Baron, and Michael A. Slote (Oxford, 1997), pp. 92-174.

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 a better outcome for others, then in most cases I would be permitted to do so. If I did, I would be acting supererogatorily.²

Self-favouring options are not the only kind.³ 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. Our last drop of water would do me more good than it would you. And yet I pass you the bottle. 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 simply acting irrationally is surely not morally wrong. Indeed, often actively harming or frustrating my own interests is morally permissible. These are all species of *self-sacrificing options*.

Let's call all of these *agent-centred options*. Now, perhaps my readiness to assent to agent-centred options illicitly imports a certain kind of political liberalism into moral theory. Or maybe it's all self-serving nonsense; really I ought to maximise value—I don't have options to act suboptimally. Or perhaps there's simply no general injunction to bring about the most good, and I should stop worrying and enjoy playing Star Wars. We should not dismiss these possibilities. But before we embrace them, there are at least three philosophically interesting questions we can ask:

² One might think that taking on a great burden for the sake of a trivial benefit to others is impermissible; even if we don't take that view, we might not describe such an action as supererogatory. Thanks to a referee here. See B. Curtis, 'The Supererogatory, the Foolish and the Morally Required', *The Journal of Value Inquiry* 15 (1981), pp. 311-18.

³ M. Slote, 'Morality and Self-Other Asymmetry', *The Journal of Philosophy* 81 (1984), pp. 179-92; 180; C. D. Broad, 'Self and Others', ed. David R Cheney (1971), pp. 1-9.

First, does commonsense morality recognise agent-centred options?⁴

Second, are agent-centred options defensible?⁵

And third, can consequentialists—people who normally think that morality requires us to maximise value—accommodate agent-centred options?⁶

In my view, the first and third of these questions have been amply covered. But, for a long time at least, the second has been neglected. More has been said about self-favouring options than about self-sacrificing ones. But their justification has been explored in most detail by someone convinced of their immorality.⁷ Though his book, *The Limits of Morality*, inspired several defences of agent-centred options against his critique,⁸ I think Kagan rightly ended those discussions by noting that, as yet, nobody had given a persuasive defence of agent-centred options that coheres with a plausible overall moral theory.⁹ The two decades since have seen modest advances.¹⁰ Most philosophers working in this area have simply tried to develop a plausible extension of consequentialism that could accommodate agent-centred options, if they were justified. This is understandable, since agent-centred 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 provide a new argument for agent-centred options. It draws directly on the fundamental deontological idea that beings with moral status are not mere sites for the

⁴ Slote, 'Self-Other Asymmetry'; S. Scheffler, *The Rejection of Consequentialism* (Oxford, 1994); S. Kagan, *The Limits of Morality* (Oxford, 1989); T. Mulgan, *The Demands of Consequentialism* (2005); T. Hurka and E. Shubert, 'Permissions to Do Less Than the Best: A Moving Band', *Oxford Studies in Normative Ethics, Volume 2*, ed. Mark Timmons (Oxford, 2012), pp. 1-27.

⁵ Scheffler, *Rejection*; S. Shiffrin, 'Moral Autonomy and Agent-Centred Options', *Analysis* 51 (1991), pp. 244-54; B. Williams, *Moral Luck: Philosophical Papers, 1973-1980* (Cambridge, 1981); P. Hurley, 'Getting Our Options Clear: A Closer Look at Agent-Centered Options', *Philosophical Studies* 78 (1995), pp. 163-88.

⁶ D. W. Portmore, *Commonsense Consequentialism: Wherein Morality Meets Rationality* (Oxford, 2011); D. Dorsey, 'The Supererogatory, and How to Accommodate It', *Utilitas* 25 (2013), pp. 355-82; Mulgan, *Demands*; D. Sobel, 'The Impotence of the Demandingness Objection', *Philosophers' Imprint* 7 (2007), pp. 1-17.

⁷ Kagan, *Limits*.

⁸ D. W. Brock, 'Defending Moral Options', *Philosophy and Phenomenological Research* 51 (1991), pp. 909-13; F. M. Kamm, 'Non-Consequentialism, the Person as an End-in-Itself, and the Significance of Status', *Philosophy and Public Affairs* 21 (1992), pp. 354-89; M. E. Bratman, 'Kagan on 'the Appeal to Cost'', *Ethics* 104 (1994), pp. 325-32; J. Waldron, 'Kagan on Requirements: Mill on Sanctions', *Ethics* 104 (1994), pp. 310-24.

⁹ S. Kagan, 'Defending Options', *Ethics* 104 (1994), pp. 333-51: 928.

¹⁰ Though see Hurley, 'Getting'; Shiffrin, 'Moral Autonomy'.

realisation of value. Other arguments for options tend to appeal either to the disvalue of a life lived in pursuit of the overall good, or to our unfortunate incapacity to be motivated by Sidgwickian impartiality. Deontologists might be sceptical about the appeal to value in the first case, and the concessionary spirit of the second. My argument draws on ideas that are more in their wheelhouse.¹¹

I do not mean, however, to pick apart the competition. In Section 2, I will use familiar critiques of those arguments to pinpoint desiderata for a case for agent-centred options. Section 3 will present the argument, Section 4 considers objections. Section 5 concludes.

2. DESIDERATA FOR AN ARGUMENT FOR AGENT-CENTRED OPTIONS

There are three broad families of argument for agent-centred options: positive, negative, and what I will call intrinsic. The positive argument is most widely endorsed: if we had always to pursue the good, then we would not only be prevented from living fulfilling lives, but we would lack a kind of integrity, since our commitments would all be potentially revisable, should they prove suboptimal. Options, by contrast, allow us to lead these fulfilling lives (that's why this is a positive argument).¹² The negative argument focuses on our shortcomings as moral agents, which prevent us from being truly impartial. If we cannot be motivated to maximise, then we are not required to do so.¹³

The intrinsic approach claims that agent-centred options are not (constitutive) means to an end, nor are they concessions to human nature. Instead, they reflect an intrinsically justified entitlement not to maximise the good, grounded in the facts that give us moral

¹¹ Fellow-travellers: D. Heyd, *Supererogation* (Cambridge, 1982); Hurley, 'Getting'.

¹² E.g. Williams, *Moral Luck*; Scheffler, *Rejection*.

¹³ S. Scheffler, *Human Morality* (Oxford, 1992). There is a quite different approach to justifying moral options (and our moral reasons more generally) that gives the concept of empathy a driving role. I lack the space to consider that approach here, but it is no doubt worth pursuing. Thanks to a referee for raising it.

status. My argument is a species of this approach.¹⁴

Some of Kagan's objections to the positive and negative arguments for options are quite narrowly targeted. We can ignore those. Others identify more general desiderata—we can extract five, in particular. An argument for options (1) must not undergenerate options, nor (2) must it overgenerate them; (3) it must generate genuine options, rather than requirements; (4) it must be well grounded in an overarching moral theory; and (5) (obviously) it must do more than simply restate, in different terms, the intuition that we have such options—it must ground them in something else. Kagan's focus throughout was on self-favouring options, but I would add a sixth: (6) an argument for agent-centred options should justify both self-favouring and self-sacrificing options.

In a little more detail: both the positive and negative arguments fall short at desideratum (1), because they are best suited only to justifying options to avoid significant sacrifices. The positive argument kicks in only when pursuing the optimal outcome will significantly affect my grounding life-projects or my integrity. But we typically think that we may suboptimally pursue our own interests even when doing so is much less costly—for example, giving up an evening every weekend or so to volunteer with a local charity would hardly undermine one's integrity or projects, and yet many people think that volunteering in this way is genuinely optional. Of course, the force of this response depends on just how demanding our positive duties are—and perhaps the right answer here is that our positive duties *are* genuinely demanding, and our agent-centred options kick in only when our integrity or projects are genuinely at threat.¹⁵ Still, many of us think we have more licence to favour our interests than that. The negative argument too has most force when great sacrifices are at stake. I cannot be motivated by a stranger's well-being to the same degree as I am by, for example, my son's. I could be motivated to sacrifice

¹⁴ See e.g. Heyd, *Supererogation*; F. M. Kamm, 'Review: Non-Consequentialism, the Person as an End-in-Itself, and the Significance of Status', *Philosophy and Public Affairs* 21 (1992), pp. 354-89; Hurley, 'Getting'; Shiffrin, 'Moral Autonomy'; M. A. Slote, *Common-Sense Morality and Consequentialism* (London, 1985).

¹⁵ Thanks to a referee here.

my life to save my son's, but not to save the stranger's. But, again, it hardly seems beyond our motivational capacities to give up one evening a fortnight to help others.

An argument for options should justify the full range of options—from the trivial to the grandiose. But (2) it must not prove too much. Kagan argued that while common sense might allow us to forego benefiting others at excessive personal cost, it prohibits harming others to avoid comparable costs.¹⁶ Can an advocate of options justify these limits? Nothing in either the positive or negative arguments does so—sometimes one can only pursue one's projects by harming others; many of us would lack the motivation to respect others' rights against harm if our own lives were at stake. Of course, one can supplement these two accounts with some explanation of the difference between doing and allowing, but an argument for options is stronger if its limits are built-in.¹⁷

Desideratum (3) is more relevant to the positive than to the negative view. If my projects and integrity can *justify* not pursuing the overall good, then why don't they generate moral *requirements*? One could here invoke Joshua Gert's distinction between 'justifying' and 'requiring' reasons—perhaps one's projects and integrity can justify, but not require.¹⁸ Indeed, Kagan recognised the category of 'noninsistent' reasons, which would have this structure.¹⁹ But what resources does the positive view of options have to motivate that category? What does introducing it do besides furnish a new way of describing the phenomenon that we are trying to explain?

Obviously any argument for agent-centred options should (4) cohere with a convincing overall moral theory, rather than simply be an awkward add-on. Kagan raises (5) as a particular problem for Slote's early version of an intrinsic view, which argued that options are grounded in the freedom to choose between a range of morally permissible

¹⁶ Contrast J. Quong, 'Killing in Self-Defense', *Ethics* 119 (2009), pp. 507-37.

¹⁷ Sobel, 'Impotence'.

¹⁸ J. Gert, 'Requiring and Justifying: Two Dimensions of Normative Strength', *Erkenntnis* 59 (2003), pp. 5 - 36.

¹⁹ Kagan, *Limits*: 378

alternatives.²⁰ Kagan argues that this simply restates the intuition that we have options, without adding to it.²¹

The last desideratum, (6), is raised by Slote.²² It is obviously not an absolute requirement. If the best approach to justifying options is piecemeal, then so be it.²³ But an argument that justifies both kinds of options can draw support from the coherence of the overall picture that it presents.

Obviously much more could be said about each of these desiderata. In particular, adherents to either the positive or the negative view might push back and argue either that they can satisfy these goals, or that they are not goals worth having. So be it. My aim here is to promote a new argument for agent-centred options, not to knock down the others.

3. MATTERING UNCONDITIONALLY

The best account of moral status entails that we have self-favouring and self-sacrificing options. I'll first introduce this account of status, then show how it entails options, before arguing that it is the right account of moral status. My approach is syncretic: it combines two elements of moral status that are often kept separate. The main challenges, then, are to show that both of these elements are necessary, and that they are compatible with one another. First, though, a note about methodology.

We can ask at least two interesting questions about moral status: what grounds moral status, and what are the implications of moral status? Both questions are very close to moral bedrock. Suppose you think that moral status is grounded in rational agency, or that it has the implications that I will say that it has. What can you appeal to, in order to justify that view? You cannot dig much deeper than you have already gone. So rather than going down further into the foundations, you have to look up at the superstructure. We justify

²⁰ Slote, *Common-Sense Morality*. Mulgan, *Demands*.

²¹ For an independent statement of this concern, see M. U. Walker, 'Autonomy or Integrity: A Reply to Slote', *Philosophical Papers* 18 (1989), pp. 253-63. For a response on Slote's behalf, see Shiffrin, 'Moral Autonomy'.

²² Slote, *Common-Sense Morality*.

²³ Walker, 'Autonomy or Integrity'.

our bedrock commitments by showing how they support plausible first-order positions, of which the target view is just one. Our aim: to develop a moral theory that best balances the virtues of fit with the intuitive ‘data’, explanatory strength, and simplicity.²⁴ This is not simply a matter of achieving coherence between our intuitions and our theory—that would exaggerate the importance of fitting the data, and underplay explanatory strength and simplicity. Our account of moral status should shed light on the substantive conclusions that it grounds, rather than just being consistent with them. Some deontologists will reject the virtue of simplicity, arguing that if the subject matter is complex, the theory should be complex too. Nonetheless, Ockham’s razor still applies: we must avoid needlessly multiplying entities. And simplicity can surely break ties when all else is equal.

On, then, to the account of moral status. Moral status is indeed most likely grounded in the distinctive properties of rational agents—the distinctively human dimension of our capacity for reason (perhaps also found in some other animals). Although I think this understanding of the *grounds* of moral status probably reinforces my account of the *implications* of moral status, I cannot hope to defend both in one paper. So I will say nothing further about the grounds of status, and concentrate instead on the implications of having full moral status.²⁵ I think there are at least two.

First, if you have moral status, then your well-being matters: your well-being can affect the deontic status of some action for you and others. It can make that action permissible or required. It can also make some permissible action supererogatory. There are other ways that one’s interests could matter. They could be valuable, or could give reasons for action. I focus on ability to affect deontic statuses because this category seems more fundamental. It

²⁴ These criteria are used by Lewis to assess putative laws of nature. But they seem equally apt to moral theorising. D. K. Lewis, *Philosophical Papers, Volume II* (New York, 1983).

²⁵ As the editor pointed out in comments, I don’t believe that moral status is binary. It can come in degrees, and non-human animals plausibly have some degree of moral status (the degree depending on the kind of animal). To avoid unnecessary complexity, throughout the rest of the paper when I refer to having moral status, I mean having *full* moral status.

is also more neutral, since we can understand it in terms of reasons or value, but sceptics about reasons- or value-based approaches to morality can understand it without them. Additionally, there are other ways to matter, besides affecting what others ought to do. Your interests might, indeed probably do, also affect what others ought to feel, and what attitudes they ought to have. My focus in this paper, however, is solely on action.

Second, if you have moral status, then *you* matter, *unconditionally*. Your capacity to affect the deontic status of your and others' actions is not wholly conditioned on facts about you that are contingently true or false. If you matter unconditionally, then you do not matter solely if and to the extent that your interests are contingently at stake in some decision. If you did not matter unconditionally, then you would not be able to affect the status of your and others' reasons for action except if and to the extent that your interests were at stake. You would be a mere site for the realisation of value; a cell in a spreadsheet whose moral significance is exhausted by your contribution to the sum at the bottom of the column. Since you *do* matter unconditionally, the degree to which your interests are at stake in a choice is only ever one part of the story about how its impacting on you affects that decision.

This is a hybrid account of moral status. The first element is utilitarian in inspiration. It expresses a basic welfarist commitment that utilitarians are right to endorse, and that others should share (I argue for this below). The second element is, in spirit at least, Kantian. It echoes the idea that beings with moral status are ends in themselves, entities with dignity, not price. It does so, I think, more successfully than other similar locutions, often used to explicate our moral status. We often say that beings with moral status matter non-instrumentally, or that they matter intrinsically.²⁶ These are entailed by mattering unconditionally, but not vice versa. If you matter non-instrumentally, then you matter without regard to the good things to which you are instrumental. But it is consistent with

²⁶ See the papers cited in R. Yetter Chappell, 'Value Receptacles', *Noûs* 49 (2015), pp. 322-32.

matter non-instrumentally that one matters only to the extent that one's well-being is at stake in some decision. The well-being is *your* well-being, after all; it is not something separate from you to which helping or not harming you is instrumental.

One's theory of well-being might also describe it as an intrinsic property of the person whose well-being it is. So, your mattering intrinsically is consistent with your mattering only if and to the extent that your well-being is affected. But if you matter unconditionally then you also matter non-instrumentally, since your mattering is not conditional on bringing about some other good. And you matter intrinsically, since *your* mattering is not conditional on anything, so *a fortiori* is not conditioned on anything extrinsic to you. But if you matter unconditionally, then your mattering cannot be reduced to the extent to which your well-being is at stake in a decision.

This hybrid account of moral status both entails that we have self-favouring and self-sacrificing options, and helps to determine their limits. Beyond mere extensional fit, it also helps explicate agent-centred options.

On the first point: if we matter unconditionally, then we must have agent-centred options. If I had to sacrifice myself whenever doing so came at a lesser cost to me than the benefit that I could realise for someone else's sake—take a cost of X units of well-being, so that you can have a benefit of $X+1$ —then I would not matter unconditionally. I would matter only insofar as the well-being realised in my life contributes to the overall sum. Since I do matter unconditionally, I cannot be required to make these kinds of marginal interpersonal tradeoffs. Hence I have at least *some* self-favouring options.²⁷

Mattering unconditionally also blocks the promotion of my own interest from

²⁷ I am not arguing that this is the *only* respect in which my mattering unconditionally makes a difference to what I and others ought to do. I am simply saying that, *necessarily*, if I had to sacrifice my interests by X in order to realise a benefit of $X+1$ for someone else, then I would not matter unconditionally—my mattering would depend only on the extent to which my interests are at stake. Thanks to the editor for pressing me here.

generating requirements, at least sometimes. If I had to promote my own interests whenever doing so was best, then my mattering would be conditioned on the extent to which some decision served my interests. If I lack any moral right to choose whether to pursue my own interests, then I do not matter unconditionally. But if I have self-sacrificing options, then, for some acts, I can determine whether that act is permissible regardless of its contingent effects on my interests. So I must have at least some self-sacrificing options.

On the second point: this account of moral status sheds light on the limits of agent-centred options. Though *I* matter unconditionally, my interests also matter, as do the interests of others. Moreover, if I matter unconditionally, then that can affect what both I and others ought to do, in a way that will sometimes limit our licence to favour or undermine our own interests. I may not harm others in the pursuit of my own interests, because that would treat *them* as mere sites for the realisation of value, whose interests could be damaged whenever doing so marginally benefits my own.²⁸ And the great good that I can do by helping others entails that I may not pursue my own interests when by bearing a relatively small burden I could advance theirs to a much greater degree. They matter unconditionally, *and* their well-being matters. Recognising this means acknowledging that sometimes I have to bear costs for their sake.

A fundamental task for moral theory, then, is to work out what makes something a reasonable cost to bear for the sake of others. This is a fraught and complex question. But recognising that others have moral status means recognising that both they and their well-being can determine the deontic status of your actions. This means, at the very least, that your agent-centred options are circumscribed by your duties of easy rescue.

As for self-sacrificing options: I think that merely frustrating one's own interests, when nobody else is affected, is always permissible. Indeed, even if others are affected, but their interests are purely 'me-regarding' (for example, think of how a film star's fans would be

²⁸ W. S. Quinn, 'Actions, Intentions, and Consequences: The Doctrine of Doing and Allowing', *Philosophical Review* 89 (1989), pp. 287-312.

upset if he committed suicide), I think I have free rein to undermine my own interests. I need not, however, insist on that point. My view is consistent with believing that, since my interests matter, and I matter unconditionally, some forms of self-sacrifice are unconscionable. If I have to respect others as mattering unconditionally, then I should also do so to myself, and this may block my licence to not merely damage my interests, but to abase myself. So it may be impermissible, for example, to enter voluntarily into slavery, or to pursue a particularly servile relationship.²⁹ Of course, the view is also broad enough to cater for those who think that such forms of self-abasement are morally permissible.³⁰

On the third point: grounding agent-centred options in moral status has additional explanatory strength. It helps us see why the problem of moral options is a problem in the first place. The principal competitors to this account of moral status endorse one but not both of its elements. And if you held one of those competing views, you would almost certainly regard the defence of agent-centred options as either trivial or (at least on first principles) impossible. Suppose you endorse only the second, more Kantian implication of having moral status. Then you would likely adopt one of two views about our duties to aid others—either, because there is no general injunction to further the interests of others, we have no such positive duties. Or else we have only those positive duties necessary to ensure preservation of one another's rational agency—the thing that makes us matter unconditionally. On the first view, self-favouring options are trivial—in the absence of positive duties to others, we are of course permitted to advance our own interests however we like (consistent with respecting our moral status and that of others). On the second view, there is also likely to be considerable leeway for self-favouring action. If all humans enjoy moral status, which is grounded in their capacity for rational agency, then arguably all we need for the preservation of that capacity is the minimal satisfaction of our basic

²⁹ J. Hampton, 'Selflessness and the Loss of Self', *Social Philosophy and Policy* 10 (1993), pp. 135-65; Hurka and Shubert, 'Permissions'.

³⁰ Thanks to a referee for pressing here.

needs.³¹ Under ordinary circumstances, it should be possible for people to ensure the basic needs of others are fulfilled without enduring excessive personal cost. So: plenty of room for options.³²

Or suppose that you endorse only the first, welfarist element of moral status. Then too you should find agent-centred options uninteresting, at least at the level of first-order moral theory (you might have some indirect way of vindicating commonsense). If I matter only to the extent that my interests are at stake in a decision, then there is no space for options at all in moral theory, whether self-favouring or self-sacrificing. The only puzzle is why people continue to endorse an obviously flawed view.

I suspect, then, that anyone who finds moral options genuinely puzzling already rejects the reduction of moral status to only one of these disjuncts, and implicitly endorses something like the hybrid account. If you don't recognise that we are also sites for the realisation of value, then there is no general requirement to promote the good, so nothing against which options are an exception. The same is true if you think that our positive duties are only minimally demanding.

If you don't also recognise that we are not *mere* sites for the realisation of value, then you will be unconcerned by marginal interpersonal tradeoffs, and moral requirements to advance one's own good. Moral options are not puzzling for true believers on either side; they are puzzling only for those who acknowledge both elements of moral status. The hybrid solution is to show that they are not mutually exclusive.

If our well-being matters, but we also matter unconditionally, then we must have at least

³¹ Of course, there could be other versions of this approach that leave very little room for self-favouring options. All I need for my argument here is that there be a version for which such options are easy to justify.

³² A reviewer notes that, in the world as it is, the imperative to help those whose basic needs are going unfulfilled could be all-consuming, leaving no room for self-favouring options. If that's right, then it is true that the Kantian may have a hard time justifying agent-centred options—indeed, perhaps they would simply deny that we have such options in the real world. In fact, however, I think that the present conditions of need are not a function of scarcity, but rather of the lack of political will on the part of the global rich. This raises the distinct issue of which kinds of duties one might have to take up the slack left by others not performing their primary duties.

some self-favouring and self-sacrificing options. But why should we endorse this conception of moral status? In particular, what are the advantages of the hybrid theory over a strictly utilitarian or a strictly Kantian alternative (where each of those alternatives denies the disjunct endorsed by the other)? And is the hybrid itself sustainable?

Suppose you endorse only the first dimension of moral status. Then, provided that you think that our well-being matters in proportion to its magnitude, and that nothing else matters, then you are committed to a crude maximising utilitarianism that foreswears not only options, but also supererogation, constraints, and considerations of distributive justice. Of course, this explains why so few contemporary consequentialists are utilitarians in this simple sense. Most either recognise other values besides human welfare, or else propose indirect forms of utilitarianism that allow us to attend to more than just the relative stakes for each party affected by a decision. These more complex views obviously sacrifice simplicity for the sake of fit with common sense. But they also sacrifice explanatory strength. As is widely recognised, indirect consequentialism tends to end up simply taking the prescriptions of commonsense morality and adding the (untested, probably untestable claim) that following these prescriptions will best serve overall well-being in the long run.³³

Adding the second element of moral status is not only simpler than shifting to indirect consequentialism, it also delivers, as I will now show, insight into the justification of constraints, the nature of supererogation, the 'separateness of persons', and of course self-favouring and self-sacrificing options. By threading a needle through each of these substantive conclusions, the hybrid account of moral status is both strengthened by and strengthens them all, as each inherits the plausibility of the others.

If we are not mere sites for the realisation of value, then there must be constraints on harming us for the sake of the greater good. If I could permissibly harm you to degree X,

³³ E. Anderson, *Value in Ethics and Economics* (London, 1993).

just for a benefit of X+I to some other person, then you would not matter unconditionally.³⁴ Instead you would matter only if and to the extent that your interests are at stake. Since you *do* matter unconditionally, you must enjoy *some* additional protection against suffering this harm, such that it can be permissible only if both the harm and the additional protection are outweighed.³⁵

The hybrid account also sheds light on supererogation. In the standard case, one acts supererogatorily by bearing disproportionate costs for the sake of realising a better outcome than one that you could have permissibly brought about. You thereby go ‘beyond the call of duty’. Hence one acts supererogatorily by not taking advantage of a self-favouring option. Indirect consequentialists might license self-favouring options like these on grounds of the long-term good done through people pursuing their own projects, or of the greater likelihood of people abiding by less demanding norms.

Sometimes, however, we can act supererogatorily in ways that are unambiguously bad from the consequentialist perspective. Since we matter unconditionally, we have a licence not to pursue our own interests at least sometimes. Since the well-being of others matters, it can affect what we ought to do. If I choose to bear some cost for others’ sake, then within bounds I may do so even if the resulting outcome is overall morally worse than a permissible alternative. This would still be supererogatory, since I have borne an unreasonable cost for the sake of realising benefits for others. This is true, for example, in the dehydration case in the introduction. At some point, however, the discrepancy between the cost to the agent and the benefits for others is so great that the agent’s self-abasement constitutes a failure of self-respect, so runs into the limits on self-sacrificing options.³⁶ Indirect consequentialists could not plausibly accommodate this kind of supererogation,

³⁴ A similar idea is defended in Quinn, ‘Doing and Allowing’; F. Woollard, ‘If This Is My Body...: A Defence of the Doctrine of Doing and Allowing’, *Pacific Philosophical Quarterly* (2013), pp. .

³⁵ Obviously, some absolutist deontologists think that this additional protection cannot be outweighed.

³⁶ Curtis, ‘Foolish’. Note that not all failures of self-respect involve a significant amount of self-harm. Thanks to X for this point.

since its net well-being results would always be suboptimal.

The hybrid account of moral status also echoes the popular idea of the ‘separateness of persons’. As Rawls argued, and many others find plausible, it is wrong to apply to a group the decision procedure appropriate for an individual deciding for himself.³⁷ My incurring some cost for the sake of a benefit that I will later enjoy is different from my imposing some cost on you to realise a benefit to me. We cannot be required to make sacrifices for the sake of marginally greater benefits to others, nor are we permitted to impose costs on some for the sake of marginally greater benefits for others.

If we matter unconditionally, then marginal interpersonal tradeoffs are neither first-personally required, nor permissible for one to impose on third parties. A gain of $X+I$ to you cannot fully compensate a cost of X to me. Nor can I permissibly impose a cost of X on you for the sake of a gain of $X+I$ to a third party. So, persons are separate in the Rawlsian sense.

But notice that while the separateness of persons might imply that it would be *impermissible* for me to regard gains to you as wholly offsetting the same sized costs to me, the hybrid account of moral status forecloses this possibility. You matter unconditionally, so if you want to sacrifice your interests for the sake of averting a same-size, or even smaller cost to someone else, then you have a licence to do so. If we matter unconditionally then we must we have self-sacrificing options; by contrast, the separateness of persons idea is silent on that point. So, the hybrid theory entails the most plausible implication of the separateness of persons but not vice versa; the intuitive plausibility of the latter idea lends support to the hybrid account of moral status, without simply restating it.

If you sign up to both disjuncts in the hybrid account of moral status, you have a unified foundation for options, constraints, supererogation, and the separateness of persons. We could adduce further support by invoking the hybrid account’s implications for egalitarian

³⁷ J. Rawls, *A Theory of Justice* (Oxford, 1999).

theories of distributive justice, as well as the moral significance of different kinds of harmful agency (in particular the opportunistic/eliminative agency distinction).³⁸ Adopting only the first implication of moral status cannot deliver such a plausible overall theory, except through some form of indirection whereby it sacrifices both simplicity and explanatory strength.

Perhaps, however, one could argue that endorsing the first disjunct leaves open the question of precisely *how* one's interests matter. Perhaps we could then distinguish between your having reasons to promote the well-being of those with moral status, and reasons to respect it.³⁹ I think that when properly fleshed out, this would be a notational variant on my hybrid approach defended above. In my terms, it would amount to saying that your well-being matters, but that you also matter unconditionally. Consider, for example, the possibility that people might have licence to give their own well-being greater weight in their deliberations than its impartially considered worth. Where does that additional weight come from, if not from recognising that something matters here besides just the quantum of well-being? What matters, if not the being with moral status herself?

But could we endorse only the second disjunct in the hybrid account? Why not simply argue that people matter unconditionally, and leave their well-being out of it? This would guarantee constraints and the separateness of persons. It would also ground options—trivially, since there would be no general requirement to advance the good of others, so we would never have to check our self-interest for the sake of others. Why should the quasi-Kantian concede any ground to the utilitarian at all?

I am baffled by the idea that our licence to favour our own interests over others' could be absolute in this way. Warren Quinn argues that if I can be required to advance the good

³⁸ W. S. Quinn, 'Actions, Intentions, and Consequences: The Doctrine of Double Effect', *Philosophy and Public Affairs* 18 (1989), pp. 334-51.

³⁹ Pettit, 'The Consequentialist Perspective'.

of others, then I am—or, more precisely, my body is—made into a means by morality.⁴⁰ But this is outlandish. Morality must provide us with some scope to pursue our own ends regardless of the overall good. But when the actual cost to me is slight in comparison with the good done, it seems outrageous to suggest that I *always* have the option to favour my own interests.

However, appeals to intuition may fall on deaf ears. Recalcitrant deontologists have heard them all before.⁴¹ Better, then, to focus not on intuitive fit, but on explanatory strength. In particular, to reject the first element of moral status deontologists have to insist on two related bright-line distinctions for which they cannot provide adequate theoretical support. The first is between the justifying and requiring dimensions of normative strength.⁴² The second is between doing and allowing harm.⁴³

Nobody can deny that my and others' well-being can justify actions that would otherwise be irrational, or indeed impermissible. Supererogation makes no sense if the well-being of others does not matter at all. And most deontologists agree that it can *sometimes* be permissible to impose severe costs on some, for the sake of a much greater benefit to many others. These verdicts are incomprehensible if the well-being of those with moral status lacks *any* capacity to affect deontic verdicts. The real question, then, is whether well-being can have this *justifying* dimension of normative strength, without also having the *requiring* dimension. Can some consideration, the magnitude of which can vary to an arbitrary degree, justify basically *anything* at the appropriate magnitude, but, regardless of magnitude, *never* generate a requirement to act?

⁴⁰ Quinn, 'Doing and Allowing'. Frances Kamm implies endorsement for the 'no positive duties' thesis in Kamm, 'Review: Non-Consequentialism, the Person as an End-in-Itself, and the Significance of Status'. For a canonical Kant-inspired view that we have no positive duties, see e.g. R. Nozick, *Anarchy, State and Utopia* (Oxford, 1974).

⁴¹ P. Singer, 'Famine, Affluence, and Morality', *Philosophy and Public Affairs* 1 (1972), pp. 229-43.

⁴² J. Gert, 'Normative Strength and the Balance of Reasons', *The Philosophical Review* 116 (2007), pp. 533-62.

⁴³ Sobel, 'Impotence'; Kagan, *Limits*.

This possibility is conceivable.⁴⁴ But is it plausible? Could one be permitted to kill one person to save a sufficiently large number of others, and yet never required to bear some much lesser cost, for the sake of realising as great a good? To believe this, one would have to believe that a consideration's weight along one dimension has *absolutely no* implications for its weight on the other dimension. I agree that we cannot just read off a reason's strength along one dimension from its strength along the other, but this is quite different from saying that a consideration (or set of considerations) could have arbitrarily high finite strength along one dimension, but *none* along the other. In other words: if there is some number n of people such that it can be permissible to kill one person to save n people's lives, then there must be some number m such that one is required to sacrifice one's life to save them. And there must certainly be some number $m-x$ such that one is required to bear some lesser cost for their sake. If this is true, then we must acknowledge that my and others' well-being can affect the deontic status of our actions, and not only by justifying what would otherwise be impermissible—they can generate genuine requirements. If this is right, then we cannot treat agent-centred options as trivial, and we need something like the present argument in their favour.

The second point is related to the first. All deontologists think that we can be required to bear very high costs to avoid harming others. I might be required to sacrifice my life, rather than inflict a lesser harm on some other person. Can one coherently endorse this view, while thinking that we can never be required to bear even relatively small costs to advance others' well-being? This seems especially implausible, since the categories of doing and allowing harm are not natural kinds with clearly defined parameters. The recent history of deontological ethics is a testament to the porousness of the boundaries between these two categories, as well as the multitude of other distinctions in modes of agency that cross-cut and overlap the doing/allowing distinction. There are clear cases of doing, and

⁴⁴ Gert defends something like this view, though he argues that the well-being of others cannot generate *rational* requirements. My focus is on moral normativity. Gert, 'Requiring and Justifying'.

clear cases of allowing, and a great mess in between, that includes, for example, enabling, redirecting, withdrawing, risking, and so on. Other things equal, I agree that where our intentions are good, we can be required to bear greater costs to avoid doing harm than to avoid allowing harm. But the doing/allowing distinction is too messy to sustain a radical normative contrast; it is a difference of degree, so cannot sustain a normative difference of kind, such as the difference between being able to require *anything* of a person, and being able to require *nothing* of her.

Deontologists cannot plausibly deny that the well-being of beings with moral status can affect the deontic status of their actions. That denial is not only counterintuitive, it also implies that the promotion of others' well-being can justify anything, but require nothing, and implies that we can be required to bear any cost to avoid a harm being connected to us in one way, but required to bear none to avoid it being connected to us in another way, when the difference between those two forms of connection is vague at best. Justifying and requiring are two distinct dimensions of normative strength, but they are not *that* distinct. Doing and allowing are indeed different, usually. But they are not *that* different.

Of course, this does leave open the prospect that one could endorse only the second element of moral status, but still think that we have positive duties to aid others insofar as that is necessitated by what makes them matter unconditionally. For example, we might think that we have such duties as enable the preservation of the rational agency of others, and nothing more than that. In the end I have little to say against this view: like mine, it would justify extensive agent-centred options, since it would retain the idea that I matter unconditionally, and so would have to include the entailments of that commitment. Like mine, it is grounded in an account of moral status. It would provide for a narrower set of duties to others, and a greater scope for the agent-centred prerogative than my account. But in essence it is a fellow-traveller.⁴⁵

⁴⁵ Thanks to a referee for helping me to see this point.

I cannot survey every competing account of the implications of moral status. And it would be unnecessary to do so if the hybrid account were fundamentally unstable. But is it? What bars us from recognising both that people's well-being matters, and that they matter unconditionally? There is no conflict between these two views. They simply mean that in every choice where a person is affected, how she matters cannot be reduced wholly to the effect of the choice on her interests (though her interests do matter, when they are at stake). The Hybrid View forfeits something in terms of simplicity to its competitors, since it identifies two implications of moral status rather than one. But this is hardly a significant cost given the benefits it yields in terms of fit and explanatory strength.

In section 2, I noted that one's argument for moral options should be more than just a restatement of the belief that we have such options (desideratum 5). The hybrid account embeds options in normative bedrock—an account of moral status. As foundations go, this account of moral status is not *much* deeper than the substantive conclusion that it underpins. And yet it does not simply restate a conviction in options. It shows, instead, that options are a necessary entailment of an attractive conception of moral status, that also entails and illuminates many other basic commitments of commonsense morality. The idea that we matter unconditionally is more fundamental than the idea that we have agent-centred options: it entails other basic deontological commitments, where an endorsement of options does not. For example, one could defend options on grounds of our inability to be motivated by the well-being of others, while denying that morality includes constraints.

Of course, from such an abstract argument, we cannot hope for any detailed verdict on precisely *which* agent-centred options we have. The argument shows only that we must have some. It does, however, satisfy desideratum 1 above, because it rules out being required to make marginal interpersonal tradeoffs even when the interests at stake are relatively trivial. It also meets desideratum 2, since the account of moral status that grounds

options also grounds constraints, which limit the exercise of those options. Desideratum 3 is likewise satisfied, since the account of moral status grounds both self-favouring and self-sacrificing options: if I matter unconditionally, then I must be morally permitted to disregard my own well-being even when doing so is suboptimal overall. I can pursue my own interest because I matter unconditionally, not only to the extent that my interests are satisfied. But if I don't want to pursue my own interest, I am not required to do so in at least some cases. Desideratum 4 is a given: this account of moral status neatly underpins a relatively uncontroversial take on commonsense morality. Desideratum 5 is satisfied, because this conception of moral status entails both that we have options, and other fundamental deontological commitments, whereas the idea that we have options does not, on its own, entail those other commitments. So we have grounded moral options in something other than just a restatement of our belief in those options. And the notion of mattering unconditionally allows us to vindicate not only self-favouring options, but also self-sacrificing options, as per desideratum 6.

4. OBJECTIONS

Four objections to my approach: have I offered a straitened menu of ways to think about status? Is my account really more attractive than the alternatives? Can one accommodate the worry about being a mere site for the realisation of value without endorsing the second implication of moral status? And last, does my account of moral status really entail options as well as constraints, rather than constraints alone? I consider each in turn.

On the first objection: I obviously haven't considered all the possible ways one might understand moral status. I have considered my hybrid approach alongside alternatives that reject one of the two disjuncts, but there are many other possibilities. If some others are independently attractive, and provide as much support as my hybrid account for the other

fixed points in commonsense morality, without supporting options, then that would undermine my argument. I cannot hope to conceive of and deflect every competing account of moral status in a single paper. It should be enough to articulate and positively argue for one, and defend it against the two most plausible neighbours.

On the second objection: Frances Kamm once defended constraints by appealing to the idea that we are ends in ourselves, so cannot be used to advance the greater good.⁴⁶ Kagan responded: on that account of status, we can't be used for the sake of others, which means we enjoy a privileged status.⁴⁷ But on his account, we can give others reasons to bear costs for our sake, and in virtue of *that* enjoy a privileged status. Which kind of privilege is better?⁴⁸

The hybrid theory of status can answer this objection: we do not have to choose. If your account of moral status denies that one can be required to bear costs for the sake of advancing another's well-being, then it is impoverished. But Kamm is also right, that if one's interests can be sacrificed just for the sake of a marginal gain to someone else, then one's moral status is to that extent diminished. It is also diminished if one is required to sacrifice one's own interests for the sake of a marginal benefit to someone else. Of course, we still need to flesh out the details of precisely which options we have. The hybrid theory of moral status is not a complete moral theory. But it can blunt Kagan's criticism, without sacrificing the virtues of Kamm's account.

On the third objection: in recent work, Richard Yetter Chappell has argued that, insofar as it constitutes a genuine, non-trivial objection to consequentialism, the 'separateness of

⁴⁶ Kamm, 'Status'.

⁴⁷ Kagan, 'Defending Options'.

⁴⁸ Compare Sobel, 'Impotence'.

persons' objection can be overcome.⁴⁹ He is interested in constraints, not options. But his objection carries over to options, as well as to the sites of value argument.

Yetter Chappell considers some implausible interpretations of the separateness of persons objection, and dismisses them. He then proposes the following gloss: a theory is vulnerable to the separateness of persons objection if, when licensing interpersonal trade-offs, it implies that the benefit achieved fully cancels out the cost, such that there are no grounds for compunction or regret, no loss to mourn. On the deprecated view, people are fungible in the way that bank notes are fungible—if you lose a tenner and gain two fivers, then you haven't lost anything at all. All that matters is the net result. Some consequentialists might think of people and their interests like this: if one person suffers a loss, but two others enjoy a benefit that aggregates to the same amount, we need only attend to the neutral net result. Only aggregate welfare is non-instrumentally valuable; people are merely instruments to that end.

This is what Yetter Chappell (drawing on Peter Singer) means by a 'value receptacle': people are to value (well-being) as banknotes are to money. This *would* be a gross mistake in a moral theory. But, as Yetter Chappell rightly observes, it is not one that utilitarians need make. Instead of treating people as fungible, and considering only aggregate well-being to be non-instrumentally valuable, they clearly can care independently about each person affected by an interpersonal trade-off, and regret the cost that someone has to bear for the sake of the overall good. They can assign non-instrumental value to each individual's well-being, rather than just the aggregate. But this has no bearing on what we ought to do—a cost of X to you is still worth only X, so can be outweighed by a benefit of X+1 to someone else.

Yetter Chappell's treatment of the separateness of persons has several weaknesses. First, what does it mean to value aggregate well-being non-instrumentally, while valuing

⁴⁹ Yetter Chappell, 'Receptacles'.

the well-being of individuals only instrumentally? The view seems a straw man. Aggregate well-being is composed of the well-being of individuals. If aggregate well-being is a non-instrumental value, then individual well-being is a non-instrumental value, since aggregate well-being just is all the individual well-being taken together. Treating different people's well-being as totally fungible is a conceptual mistake, hence not a charitable interpretation of the separateness of persons objection.

Second, Yetter Chappell thinks the objection has to do only with attitudes. His token-pluralistic utilitarianism can, in its deontic verdicts, be extensionally identical to token-monistic utilitarianism (according to which only aggregate well-being is non-instrumentally valuable), but preferable since it encourages us to adopt the appropriate attitude to the losses inflicted in the pursuit of the overall good. This misunderstands the separateness of persons worry. It has nothing to do with our attitudes: it concerns instead what we ought to do. We ought not assume that benefits to one person can cancel out same-sized costs to another.

To see that attitudes are beside the point, notice that it's often bad to treat costs and benefits as fungible, even when it is acceptable to treat an entity as a mere site for the realisation of value. Suppose, for example, you are managing a population of kangaroos in a nature reserve. It might be necessary to cull some to allow the remainder to flourish—perhaps overpopulation is causing starvation and disease. But it would be extraordinarily callous to feel no ambivalence about such an extreme endeavour. The deaths of the culled 'roos are not 'cancelled out' by the good realised by the overall population. You shouldn't care only about the net result. And yet these kinds of marginal intraspecies trade-offs are completely defensible, since kangaroos do not have the same kind of moral status as do humans, so treating them as mere sites for the realisation of value is acceptable. We may justifiably ignore the separateness of kangaroos.

Indeed, even when making intrapersonal trade-offs, it is often a mistake to regard costs

and benefits as cancelling one another out. If I endure hardship now for the sake of a benefit later—getting up at all hours to take Moab, my Labrador, outside so that he'll be housetrained—then the cost and benefit are not fungible with one another. I suffer a cost now—a deficit in non-instrumentally valuable well-being—for the sake of a later benefit. The benefit is great enough to justify the cost, but does not cancel it out. Yetter Chappell offers the 'fungibility' worry as an interpretation of the separateness of persons objection. But it is as much a mistake to consider intrapersonal costs and benefits fungible as it is to regard interpersonal costs and benefits in this way. Since the point of the objection is the difference between intrapersonal and interpersonal trade-offs, this cannot be an adequate interpretation of that objection.

Here is a more promising interpretation. Valuing Gareth or Geraint's well-being non-instrumentally is not the same as valuing Gareth or Geraint unconditionally. Yetter Chappell's token-pluralistic utilitarian still cares about Gareth or Geraint only insofar as their well-being makes the world a better place. They are still mere value receptacles, even if their well-being is not fungible. To value them unconditionally, you have to recognise that each can determine what you ought to do, independent of contingent facts about how your action affects his well-being. Otherwise what you value is the well-being, not Gareth or Geraint, and you're disrespecting them by treating their lives as containers for the stuff that you really do care about. And since you care only about the value realised in their lives, rather than about Gareth or Geraint, where the value is located makes no difference to what you ought to do. So you can sacrifice Geraint's interests for a marginally better outcome for Gareth. Of course, you care about well-being non-instrumentally, so you recognise that Geraint suffers a genuine loss. But since he matters only contingently, insofar as your actions affect his well-being, this loss doesn't affect what you ought to do. You ignore the distinction between persons, because you don't think that each *person* is unconditionally valuable, separately from the well-being realised in his life.

Of course, this presupposes that we *should* value Gareth and Geraint unconditionally, rather than caring only about their well-being. In other words, it presupposes my account of moral status. And a utilitarian would reject that account. The force of the separateness of persons objection then reduces to the force of the argument for this account of moral status. Does this make the objection trivial, in Yetter Chappell's sense? I do not know. It certainly makes it external, in a way that Yetter Chappell's version is not. But that strikes me as exactly right. The separateness of persons objection, and the deeper sites of value idea that underpins it, pinpoint a fault-line between ethical theories. The disagreement here is profound. We should not expect it to be translatable into an internal objection, grounded in shared premises. At this level, all putative foundations are controversial. We can do no more than make the best case we can for our different understandings of the foundations of ethics.

The preceding objections questioned either my menu of conceptions of moral status, or my choice from that menu. But one might think that we can make the same choice that I made, without endorsing agent-centred options. Could one believe both that we are sites for the realisation of value, and that we are not mere sites for the realisation of value, and yet deny that we have options?

The simplest route to this conclusion would be to endorse constraints, but reject options. Then you might think that our status as ends in ourselves is properly acknowledged, because other people may not *treat* us as mere sites for the realisation of value. But we ourselves *are* required to sacrifice our interests whenever doing so is optimal.

I think this is a mistake. How could one believe that we matter unconditionally, and yet believe that it is wrong to sub-optimally sacrifice our own interests, and mandatory to bear costs for the sake of a net overall improvement? It is inconsistent to justify constraints by appeal to the sites idea, and then deny that it has these further implications. If you *are* a

mere site for the realisation of value, since you can be required to bear a cost of X for the sake of a benefit of X+I to someone else, then it should be permissible for others to *treat you* as such a site, and impose that cost on you for the other person's benefit. We are prohibited from treating others as mere means only if they are not mere means!

Of course, one could counter here that sometimes we have a 'right to do wrong'—it might be wrong for others to force me to bear a cost that I am required to take on myself. But in the plausible examples of this phenomenon, some specific factor explains why others may not interfere. For example, a father might be morally required to love his children. But it would be inappropriate for anyone else to enforce that requirement because doing so would be self-defeating—love is probably not in our voluntary control, and even if it was, would lose value if the father was forced to love. It might also be morally wrong to form some kinds of thoughts, but others couldn't enforce our obligations in this case, because they lack access to our minds. But it is wholly implausible that the full range of agent-centred options, as well as supererogatory action, can be explained in this way. The sites idea entails that we have options even when, if we didn't have that option, it would clearly be permissible for others to enforce our moral obligations. The most obvious examples of this are monetary: I have a moral option to waste my own resources if I want to, and to forebear from making sacrifices for the benefit of others beyond a given point, because I am not a mere site for the realisation of value. If I was morally required to bear those costs for others' sake, then it would obviously be permissible for others to enforce those moral requirements, for example through taxation. The right to do wrong neither adequately explains our options to act suboptimally, nor successfully captures the idea that we are not mere sites for the realisation of value. It is a red herring.

5. CONCLUSION

If we were required to sacrifice our own interests whenever doing so was best overall, or prohibited from doing so unless it was optimal, then we would be mere sites for the realisation of value. Our interests, not ourselves, would wholly determine what we ought to do. We are not mere sites for the realisation of value—instead we, ourselves, matter unconditionally. So we have options to act suboptimally. These options have limits, grounded in the very same considerations. Though not *merely* such sites, you and I are *also* sites for the realisation of value, and our interests (and ourselves) must therefore sometimes determine what others ought to do, in particular requiring them to bear reasonable costs for our sake. Likewise, just as my moral status grounds a requirement that others show me appropriate respect, so must I do to myself. I suggested above that mere self-harm (damaging my own interests) is always permissible, but some forms of self-abasement are not. Much more needs to be said; in particular, while the present argument ensures that we have *some* self-favouring and self-sacrificing options, it leaves open their precise contours. I leave those details for another time; my aim here has been to show that agent-centred options are directly grounded in an understanding of moral status that deontologists can adopt, and which we should adopt, because it balances the virtues of simplicity, strength and fit better than its competitors.⁵⁰

seth.lazar@anu.edu.au

⁵⁰ I presented earlier versions of this paper at the Australasian Moral Philosophy Workshop, at Kioloa, and at seminars in Yale and Berkeley. My thanks to the organisers of those meetings for inviting me, and to the participants for their many helpful comments. For comments on drafts, thanks to Garrett Cullity, Meir Dan-Cohen, Dale Dorsey, Tom Hurka, Shelly Kagan, Niko Kolodny, Chris Kutz, Veronique Munoz-Darde, Shmulik Nili, Doug Portmore, Nic Southwood, and Yuan Yuan. Thanks also to a reviewer and the editor of this journal, for their many helpful and insightful suggestions. This paper was supported by Australian Research Council grant DP170101394.