

SELF-OWNERSHIP AND AGENT-CENTERED OPTIONS

BY SETH LAZAR

Abstract: I argue that agent-centered options to favor and sacrifice one's own interests are grounded in a particular aspect of self-ownership. Because you own your interests, you are entitled to a say over how they are used. That is, whether those interests count for or against some action is, at least in part, to be determined by your choice. This is not the only plausible argument for agent-centered options. But it has some virtues that other arguments lack.

KEY WORDS : interests, agent-centered options, well-being, reasons for action

I. INTRODUCTION

Here is a simple view of how morality works. You ought always to perform your best option. This is a “hard” ought. If there’s a better alternative available, and you choose a worse one, then you have done the wrong thing. And what, after all, could one say against such a view? If you know all the relevant facts, and all the reasons have been accounted for, then what rational basis can you give for *not* performing your best option?

And yet, as compelling as it is, many of us reject this view of morality. We may sometimes, *without wrongdoing*, forego our best option. Sometimes we can serve our own good when, strictly speaking, the best option would be to engage in self-sacrifice. And sometimes we may sacrifice our own good, even though our best option would be to serve it. These are *self-favoring* and *self-sacrificing* options. Together, we can call them agent-centered options. Though they are widely endorsed, there is no good justification for them. This essay aims to provide one.

Here is my central thesis. Everyone’s well-being can make the world a better place. Your well-being does too. That can give everyone—including you—reasons for action. But your well-being is instantiated in *your* life. It makes the world better only through making *your life* better. When some fact that is instantiated in your life has reason-giving force, you should have some say over how that fact features in the justification of your actions and the actions of others. I describe this as a species of self-ownership. To be a self-owner is to have a special authority over one’s own life. Since your well-being is part of your life, you should have a special authority over your well-being. If you lack agent-centered options, then you have no such authority. So you must have agent-centered options.

In Section II, I set out some desiderata for my argument to satisfy. I expand on the argument in Section III. Section IV considers objections. Section V concludes.

II. EXISTING ARGUMENTS FOR AGENT-CENTERED OPTIONS¹

Arguments for options come in three broad categories: positive, negative, and what I will call intrinsic. The positive argument is the most popular. It holds that if we had to always realize the impersonally best outcomes, then we would be unable to live fulfilling lives; we would lack integrity, because all our “commitments” would be in principle revisable, should they prove suboptimal. Options allow us to flourish and live with integrity (hence the name “positive”).² Negative arguments emphasize our shortcomings as moral agents. According to this type of argument, we could never be motivated to be genuinely impartial, and moral theory must defer to our motivational capacities. So, since we cannot be motivated to maximize, we are not required to do so.³

The intrinsic approach argues that agent-centered options are neither (constitutive) means to an end, nor concessions to our weaknesses. Instead, they are grounded in an antecedent intrinsically justified entitlement.⁴ In this essay, I put forward one such ground.

The most extensive treatment of agent-centered options concludes that the negative and positive arguments fail, as do early versions of the intrinsic approach.⁵ Some of Shelly Kagan’s objections to the positive and negative arguments are narrowly targeted against those arguments, so can be ignored here. Others are more general, and can be used to pick out some pitfalls that any argument for agent-centered options should avoid.

First, we must avoid undergenerating options. Both the positive and negative arguments fall short here, because they are best suited only to

¹ In this section, I cover the same ground as I do in Section 2 of “Moral Status and Agent-Centered Options,” *Utilitas* 31, no. 1 (2019), 83–105.

² There are too many proponents of this view to cite them all, but the most influential sources are Bernard Williams, *Moral Luck: Philosophical Papers, 1973–1980* (Cambridge: Cambridge University Press, 1981); Samuel Scheffler, *The Rejection of Consequentialism* (Oxford: Clarendon Press, 1994).

³ See especially Samuel Scheffler, *Human Morality* (New York: Oxford University Press, 1992).

⁴ The only other examples I have found along the same lines are David Heyd, *Supererogation* (Cambridge: Cambridge University Press, 1982); Paul Hurlley, “Getting Our Options Clear: A Closer Look at Agent-Centered Options,” *Philosophical Studies* 78, no. 2 (1995): 163–88; Seana Shiffrin, “Moral Autonomy and Agent-Centered Options,” *Analysis* 51, no. 4 (1991): 244–54; Michael A. Slote, *Common-Sense Morality and Consequentialism* (London: Routledge, 1985). Note that Garrett Cullity’s argument is rather hard to place: he defends options on grounds of consistency, arguing that we must be permitted to pursue our own projects, since we are permitted to help others pursue theirs. Garrett Cullity, *The Moral Demands of Affluence* (New York: Oxford University Press, 2004).

⁵ Shelly Kagan, *The Limits of Morality* (Oxford: Clarendon Press, 1989).

justifying options that involve significant sacrifices to the agent.⁶ The positive argument does not kick in unless my projects or integrity are at stake. But self-favoring options are often much less dramatic than this. Volunteering one evening a fortnight with a local charity would hardly undermine one's integrity or projects, and yet we genuinely think of it as volunteering, rather than doing one's duty. The negative argument is also most plausible for significant sacrifices. Though I might not be able to bring myself to sacrifice my life to save a stranger's, it is hardly beyond my motivational capacities to give up one evening a fortnight.

Just as an argument for agent-centered options should avoid false negatives, it must also be wary of false positives. Though we can forego self-sacrifice that would benefit others, common sense suggests that harming others to benefit oneself is very hard to justify.⁷ An advocate of options must explain their limits—all the better if the same considerations that justify options also limit them. Both the positive and negative arguments come unstuck here. Just as pursuit of one's projects can necessitate missing opportunities to help others, it can also necessitate harming them. And when our vital interests are at stake, many would find it too hard to obey a prohibition on harming others to benefit oneself. Why, then, don't we have these kinds of options?

The third pitfall is to fail to justify genuine options, rather than requirements. If I can be justified in foregoing the overall good for the sake of my projects and integrity, then why am I not required to do so? After all, if the balance of reasons—taking projects, integrity, and so on into account—supports favoring my own interests, then what could justify me in not pursuing my own good? Adherents to the positive view could invoke Gert's distinction between "justifying" and "requiring" reasons—perhaps it's a brute fact about one's projects and integrity that they can justify, but not require.⁸ But Kagan already recognized the category of "noninsistent" reasons.⁹ His complaint was that it does nothing more than redescribe the phenomenon that we are trying to explain. The negative argument, note, avoids this objection.

Kagan focused on the positive and negative views; the only intrinsic view on the table at the time was Michael Slote's. Kagan's counterargument was simple and compelling. Slote argued that options are grounded in the importance of moral autonomy—that is, freedom to choose from a range of morally permissible alternatives.¹⁰ Kagan objected that this simply restates the intuition that we have options,

⁶ Kagan, *Limits of Morality*.

⁷ Not everyone agrees on this point, e.g. Jonathan Quong, "Killing in Self-Defense," *Ethics* 119, no. 3 (2009): 507–537.

⁸ Joshua Gert, "Requiring and Justifying: Two Dimensions of Normative Strength," *Erkenntnis* 59, no. 1 (2003): 5–36.

⁹ Kagan, *Limits of Morality*, 378.

¹⁰ Slote, *Common-Sense Morality*.

without actually justifying it.¹¹ Views that justify agent-centered options intrinsically, rather than by appealing to what they are good for, are particularly vulnerable to this challenge.

The last consideration worth mentioning was not on Kagan's radar, because he focused almost exclusively on self-favoring options. I agree with Slote that a defense of agent-centered options is stronger if it can offer unified foundations to self-favoring and self-sacrificing options.¹² This is obviously not a *sine qua non*. If the best approach to justifying options is piecemeal, then so be it.¹³ But simplicity is surely a virtue in a moral theory, as is coherence. An argument that successfully justifies both kinds of options can draw support from the coherence of the overall picture that it presents.

Obviously much more could be said here—adherents to the negative and positive views will undoubtedly have answers to some of Kagan's objections. But my concern is only to pick out some desiderata for my argument to meet, rather than to comprehensively assess the merits of the alternatives. And the argument will be stronger if it neither undergenerates nor overgenerates options, if it grounds genuine options, rather than requirements, if it is properly motivated, and if it explains both self-favoring and self-sacrificing options on the same grounds. I think my scorecard is four out of five. The view might undergenerate options. But it meets the other desiderata. I will come back to this in the conclusion.

III. IT'S MY WELL-BEING!

The argument begins with a controversial bedrock commitment. I take it to be self-evident. If you reject it, then the rest of the argument will not speak to you.

Here is that bedrock: my well-being, just like yours, and anyone's, can make the world a better place.¹⁴ And we have reason to make the world a better place. The mere fact that some act benefits either of us gives us reason to perform it. And sometimes—perhaps often—it is wrong not to make the world a better place, when we have the opportunity to do so.

More needs to be said about what "well-being" means here. I aim to be neutral between different theories of well-being, up to a point. For on

¹¹ For an independent statement of this concern, see Margaret Urban Walker, "Autonomy or Integrity: A Reply to Slote," *Philosophical Papers* 18, no. 3 (1989), 253–63. For a response on Slote's behalf, see Shiffrin, "Moral Autonomy."

¹² Slote, *Common-Sense Morality*.

¹³ Walker, "Autonomy or Integrity."

¹⁴ Perhaps not only human well-being; since nonhuman animals lack the rational capacities either to sacrifice their interests for the greater good, or to withdraw their interests from consideration in calculating the greater good, I will focus only on human well-being in this essay.

some theories of well-being my bedrock commitment is clearly false. If you think it advances a sadist's well-being to inflict pain on a puppy, then clearly *that* kind of well-being doesn't make the world a better place. My own view is that satisfying sadistic preferences does not in fact contribute to well-being. If you disagree, then confine your attention to all and only those aspects of well-being that *do* make the world a better place. This should be the rule, not the exception. Let's call an aspect of your well-being an *interest*. Your interests don't have to be fully virtuous in order to make the world a better place. Innocent pleasure is enough. From now on, whenever I mention well-being, interests, or making people better off, I am referring only to those aspects of individual well-being that conduce to making the world a better place.

Plenty of people would still reject this starting point. I can't see how one could deny that making people better off in the right way makes the world a better place. But one might well deny that one has a standing duty to make the world a better place.¹⁵ On this view, agent-centered options are trivially easy to justify. Agent-centered options are puzzling only for those of us persuaded that we have a standing obligation to make the world a better place, to which agent-centered options seem an exception.

When thinking about how to make the world a better place, and taking well-being into account, we must not treat well-being as an abstract object instantiated in the world as such, but instead as something that is realized only in the lives of individuals.¹⁶ With respect to well-being, the world is a better place only insofar as individuals are made better off. Necessarily, my well-being, which makes the world a better place, is instantiated in my life. Because of this fact, I must have some measure of authority over the role that my well-being plays in justifying my actions and the actions of others. I'll show in a moment that if we grant this claim, then we must also grant agent-centered options.

But why should we grant this claim? Here is the general principle on which it rests. I have a special authority over things that are uniquely part of me. I take this to be the truth at the heart of the metaphor of "self-ownership." Why metaphor? Because ownership is an institutional normative relation, which involves many different rights, only some of which are relevant to self-ownership. It is an open question whether it even makes sense to talk about owning oneself. But at the heart of this metaphor is a very simple idea. To own something is just to have a special authority over it—minimally, for your say over what should be done with that thing to have more weight than that of others who do not own it.

¹⁵ E.g., Judith Jarvis Thomson, *Rights, Restitution, and Risk: Essays in Moral Theory* (Cambridge, MA: Harvard University Press, 1986); Warren S. Quinn, "Actions, Intentions, and Consequences: The Doctrine of Doing and Allowing," *Philosophical Review* 98, no. 3 (1989): 287–312.

¹⁶ Richard Yetter Chappell, "Value Receptacles," *Noûs* 49, no. 2 (2015), 322–32.

The fundamental idea—the truth—at the heart of self-ownership, is that if something is part of my self, then I should have special authority over it.

The next task in unpacking the idea of self-ownership is to ask what counts as the self. There are some challenging metaphysical questions to answer, some of which are taken up by other essays in this issue. Happily, however, we can sidestep most of them here. It is analytically true that one's well-being is part of the relevant self. Well-being consists, necessarily, in the flourishing of some living entity. A given quantum of well-being in the world must be instantiated in some individual's life. It makes the world a better place *only in virtue of* how it is instantiated in that life. This is the simple idea that I mean to invoke by appealing to the "self." The key point is not to offer a coherent account of what the self is, but to identify the underlying normative properties at stake. Something is part of your self, in the relevant sense, if it comes into the world, and persists in the world, *only through you*. If you were to be suddenly removed from the world what, *necessarily*, would go with you? I'll come back to the "necessarily."

Let's grant, then, that well-being makes the world a better place only in virtue of making someone's life better. Why should that lead us to think that each person has a special authority over her own well-being? The underlying idea is very simple. If I bring something into the world—if it is in the world, solely and necessarily, in virtue of how it is instantiated in my life—then I must have some discretion over that thing. We have little trouble understanding the similar idea that, if I am solely responsible for creating some object—a sculpture, say—I should have special authority over that object. The mere fact that this object would not exist but for you gives you some say over it. If there are others who stand in the same kind of relation to the sculpture, then they may also have some say. Those who are wholly unrelated to the sculpture have no such say. Your well-being is uniquely a part of your life. If you did not exist then, necessarily, your well-being would likewise disappear. So you have a say over it that others, whose metaphysical relation to your well-being is more remote, do not.

The concept of self-ownership is usually used to place limits on what others may do to me—in virtue of my being a self-owner, others may not impose costs on me just for the sake of realizing a net improvement in the world as a whole. In other words, self-ownership is normally deployed to justify *constraints*. But in my view, constraints and options are two sides of the same coin. Just as we may not impose costs on you for the sake of a marginal social improvement, nor can you be morally required to bear such costs for the sake of a marginal improvement. The underlying rationale is the same.

I do not claim, of course, that our ownership over our well-being exhausts the significance of self-ownership. We plausibly have just the same kind of authority over our bodies and thoughts, generating rights to bodily integrity and to freedom of thought. What's more, that authority

most likely strengthens the case for agent-centered options.¹⁷ But the underlying idea in each of these cases is the same: if some property exists in the world only in virtue of how it constitutes my life, then whatever capacity that property has to give me and others reasons for action must be to some extent in my power to control. Indeed, it is intolerable to suppose that an individual might have no more say over the constituent elements of her own life than does the moral community at large.

One might question, at this point, how my view will make sense of other-regarding interests. I care about my son's flourishing. When his life goes better, my life goes better. So his well-being makes the world a better place not only by being realized in his life, but also through its effects on me. Does this mean that I should get the same kind of say over my son's interests as I have over my own? It does not. First, the benefit to me arises only *in virtue of* the benefit to him. His well-being is still the ground of this improvement. That well-being is instantiated in *his* life. Second, my son's flourishing *necessarily* makes the world a better place by instantiating a given quantum of well-being in his life. Contingently, it also causes some additional well-being, which is instantiated in my life. I have a say over the well-being instantiated in my life. But not (or at least not on these grounds) over that instantiated in my son's life.

The next step in the argument is to show that, for me to have a say over my well-being, I must have agent-centered options. Let's return to the analogy of ownership. I suggest, following Warren Quinn, that we cannot say that one owns X, if it is true that what one does with X is wholly constrained by one's moral reasons.¹⁸ To own X, you have to have some measure of authority over it. This means that you get to use it without being subject to coercion or moral criticism by others. But if morality is maximally prescriptive, such that you are always required to use X to perform your best option, then there is no space for you to have any authority with respect to X. Perhaps sometimes there will be different ways of realizing the same outcome. But in principle there need not be. If you always have to perform your best option, then you don't have any scope for asserting your own will over X. So you don't own X. As a result, if we think it important that you *should* have special authority over X, then we have to reject the idea that morality is maximally prescriptive. Minimally, we need to give you the authority to destroy X should you wish to. Plausibly, we should also recognize that you are not required to give up X just in case doing so would realize a somewhat better outcome. In other words, if you own X, then you must have agent-centered options with respect to X.

Your well-being makes the world a better place. It is necessarily and uniquely instantiated in your life. You are the locus of that well-being.

¹⁷ Quinn, "Doing and Allowing"; Fiona Woollard, "If This Is My Body . . . : A Defense of the Doctrine of Doing and Allowing," *Pacific Philosophical Quarterly* 94, no. 3 (2013): 315–41.

¹⁸ Quinn, "Doing and Allowing." See also Woollard, "If This Is My Body."

You should have a say over things that are necessarily and uniquely instantiated in your life. If you are required to sacrifice your well-being whenever doing so realizes a net improvement overall, or required to advance your well-being whenever doing so is for the best, then you would have no such say. So, you must have self-favoring and self-sacrificing options, respectively.

The case for self-sacrificing options is, I think, strongest. Suppose that you decide not to advance your interests, when doing so would have been your best option overall. What objection could really be raised against you? That you lowered aggregate well-being? But you did so only in virtue of lowering your own well-being. And surely we care about aggregate well-being only insofar as we care about the people whose well-being makes it up. What if your undermining your interests upsets a just distribution? Again, insofar as your doing so affects only you, if you have decided to bear this cost, who can gainsay you?

It's worth pausing to emphasize this point. If the fact that my well-being is instantiated in my life gives me a right to a say over how it features in the justification of my actions and those of others, such that I can permissibly frustrate my own interests at the cost of the overall good, then adherents to some more familiar arguments for self-favoring options have an answer to a prominent objection against them.¹⁹ Williams, Scheffler, and others, have argued that self-favoring options are necessary in order to develop the projects and commitments that are constitutive of a life well lived. Kagan responded: if those reasons are so powerful, such that they can make it permissible to ignore the weighty interests of others, then why aren't we morally *required* to pursue these projects and commitments? Why isn't it morally *wrong* to sacrifice oneself for the sake of realizing a better outcome overall? My argument offers an answer to this objection. The pursuit of worthwhile projects and commitments *is* central to realizing a good life. But insofar as the reason-giving force of such projects and commitments comes down to their contribution to the well-being of the agent (which undoubtedly does not exhaust their normative significance), the agent should have authority over her own well-being, which means she can undermine it (at least up to a point) without wrongdoing. That some self-favoring action is justified by these reasons does not entail that it is required, because one can faultlessly fail to advance one's own well-being.²⁰

However, I think that my argument can justify self-favoring options without appeal to Williams and Scheffler's positive arguments. If I have to sacrifice my well-being whenever doing so realizes a greater good for someone else, then I have no say over how my well-being features in

¹⁹ Williams, *Moral Luck*.

²⁰ Of course, sometimes these projects will be impersonally valuable, and *will* generate requirements.

the justification of my actions. The situation would be identical to one in which I have to pursue my well-being whenever doing so realizes the most good. In either case, what I am required to do is fully determined by the advancement of aggregate well-being. The moral landscape would be identical if I had no say over my well-being.

IV. OBJECTIONS

The argument is quite simple. Your well-being makes the world a better place. That can give you and others reasons for action, determining what it is permissible to do. But your well-being makes the world a better place only in virtue of how it is instantiated in your life. Because of that fact, you must have some measure of authority over how your well-being features in the justification of your actions and those of others. If you don't have self-favoring and self-sacrificing options, then you don't have that authority. So you must have those options.

The most natural line of response to the argument is to question the claim that, if one doesn't have self-favoring and self-sacrificing options, then one doesn't have the appropriate authority over one's own well-being. This can take two forms: one says that I can have that authority without having any such options; the other says that my having self-sacrificing options is enough for me to have that authority. I'll discuss each in turn.

On my account, I have authority over X only if having that authority makes a difference to what I am morally permitted to do. But we can understand authority differently from this. One could instead argue that I have authority over X if my doing so makes a difference to what others are permitted to do. Specifically: I have authority over X if others are not permitted to *force me* to sacrifice or promote X in some range of cases. My argument might have shown that some kinds of moral verdicts may not permissibly be enforced by others. In other words, it might have shown that we have a right to do wrong. But it has not shown that we have genuine agent-centered options.

The objection fails, because my explanation of what it is to have authority over one's well-being is more plausible than the proposed alternative. First, it seems very likely that the right to do wrong has its proper place in political philosophy, not normative ethics. It concerns, primarily, the need to restrict the power of political institutions to interfere with people's lives, in the pursuit of enforcing a particular moral perspective. It is grounded primarily in the pragmatic realization that our institutions are inevitably imperfect, and any attempt by them to enforce our general moral obligations would most likely be disastrous. It also draws considerable support from the fact of reasonable disagreement about moral questions. If this is what we mean by the right to do wrong, then obviously my argument is irrelevant to it, and there is no objection here.

But let's grant that there might be some version of the right to do wrong at the level of fundamental moral theory. Our question, then, is this: my argument says that we should have some authority over how our well-being features in the justification of our and others' actions. Is this adequately accommodated, without agent-centered options, by some version of a right to do wrong?

Much turns, here, on how we understand the right to do wrong. Where I say that it can be morally permissible to preserve your own interests, even if sacrificing them would lead to a net improvement overall, I am asserting that you have an agent-centered option to favor yourself in this case. My opponent would say that it is morally wrong to favor yourself, but that others are not permitted to force you to perform your best option. Everything hangs, then, on how we understand "force."

Let's consider two possibilities. First, if X is wrong, but I nonetheless have a right to do it, that might mean that others are not permitted to use the threat of harm, or indeed actual harm, to compel me to do X. If that's how to understand the right to do wrong, then I think it is more the rule than the exception at the level of moral theory. We are not, in general, permitted to threaten others with harm in order to get them to do what they are morally required to do. We may do so only when their wrongdoing is likely to generate costs for others, and the harm that we threaten is necessary and proportionate. On this account of the right to do wrong, it cannot adequately accommodate the special authority we have over our own interests, because there is nothing special about it being wrong for others to threaten us with harm in order to get us to do the right thing.

Moreover, enforcement needn't only involve actual force. Our most powerful means of enforcing conformity with norms is through moral censure. If I know that it is morally wrong to refrain from sacrificing myself for the greater good, and that I will be subject to moral censure if I don't sacrifice myself, then I don't have any meaningful authority over this choice. So if the right to do wrong is just a right not to be threatened with preventive harm as a means to get one to do the right thing, then it is not sufficient for having the level of say over one's interests necessitated by my argument.

Suppose, second, that we construe the right to do wrong as the right to act impermissibly without moral censure. Could this more successfully substitute for my claim that we have agent-centered options? I don't think it could. First, I find this conception of the right to do wrong practically unintelligible. If some action is wrong, then unless the agent is fully excused, I cannot see why it should be inappropriate to censure them. Think of any of the standard cases in which one might think that one has a right to do wrong—marital infidelity, failures of parental love, failures of politeness and good will, for example—in each of these cases moral criticism remains perfectly appropriate.

What's more, even if others are not allowed to blame me (even if I am blameworthy), if it is wrong for me to X, then I should feel guilty if I X. And if X-ing comes at the cost of feeling guilty for acknowledged wrongdoing, then I do not have genuine authority with respect to X.

The right to do wrong cannot substitute in my argument for the possession of agent-centered options. What of the idea that my argument vindicates only self-sacrificing options, not self-favoring ones?

My first response is to revert to the Williams/Scheffler positive defense of self-favoring options. Even if all I have shown is that we can't be required to pursue our own interests, that is still an important result, since together with the positive arguments for personal prerogatives it does entail that we have the full suite of agent-centered options, and offers a response to a prominent objection by Kagan.²¹

But I don't have to be this concessive. If I had to sacrifice my well-being whenever doing so realized a somewhat greater good for others, then I wouldn't have any say over how my well-being features in the justification of my actions and those of others.

What's more, the case for self-favoring options is identical to the case for self-sacrificing options. So if you find it convincing for one, you should also find it convincing for the other. The key point is not: they're my interests, so I am permitted to thwart them if I want to; instead it is: they're my interests, so I'm entitled to some degree of authority over them. The latter claim entails (assuming the failure of the right to do wrong objection) both that I may sacrifice them if I wish, and that I may withhold them even when they are outweighed by others' interests.

Suppose that you're convinced, at this point, that my argument does ground both self-favoring and self-sacrificing options. We might not yet be home and hosed. One might object, for example, that though my argument can vindicate *some* such options, it cannot justify all the agent-centered options that we plausibly have. For all I have said so far, perhaps we may favor our own interests only when doing so comes at a very small cost to overall well-being; perhaps we may undermine our own interests only to a very trivial degree.

This objection is not compelling for self-sacrificing options. If the advancement of my interests makes the world a better place, it does so only by affecting me, and my having authority over my own interests is really satisfied only if I can entirely withdraw them from consideration. Note, though, that this is consistent with there being some self-regarding duties. They would not, however, be grounded only in the importance of advancing one's own well-being, but in some other moral consideration.

²¹ Kagan, *Limits of Morality*.

However, the objection hits home for self-favoring options. I have indeed shown only that there must be some such options. My argument does not specify which options we have. But I stand by this less ambitious result. I think we must be pluralist about what justifies agent-centered options. To work out precisely which self-favoring options we have, we need to invoke all of the arguments in favor of those options. The procedure would be to work out how much cost it is reasonable for the agent to be required to bear, in order to bring about a given degree of moral advancement. This means considering, in particular, the positive case for options as developed by Williams, Scheffler and others, as well as a moral status-based argument for options that I advance elsewhere, and the considerations of self-ownership introduced here. Undoubtedly other such arguments would be relevant too. Only when all the reasons in favor of having options are on the table can we properly balance the good achieved against the costs borne. And even then, it would not simply be a matter of applying a continuous ratio. Many other facts are relevant to which costs count as reasonable (for example, whether the agent was responsible for this choice situation coming about).

So, it is hardly surprising that we cannot read off from the present argument a precise list of which options we have. It shows only that we have some such options. To achieve more concrete results, we need to take all the arguments for options into account, as well as the other considerations that bear on whether you can be required to bear a given cost. But we should not underestimate the distinct contribution of this argument—the positive argument notoriously fails to explain self-sacrificing options and trivial options, and risks making the pursuit of one's own interests required, rather than merely optional. The argument from self-ownership corrects each of these defects. It obviously explains self-sacrificing options and grounds genuine options, not requirements. And it can support trivial options: marginal interpersonal tradeoffs cannot be morally required even when the interests at stake are trivial.

One might also charge that my argument does not explain the limits of agent-centered options—the points at which it is no longer permissible to favor your own interests, or perhaps even to thwart them. I agree that there are some such bounds. I may not favor my own interests when sacrificing them would realize a great enough good for others; it is also plausible that there are *some* limits on permissible forms of self-harm.²²

Here I think we are on firmer ground. The argument from self-ownership supports options, but it also supports constraints, as Warren Quinn and

²² For example, Jean Hampton, "Selflessness and the Loss of Self," *Social Philosophy and Policy* 10, no. 1 (1993), 135–65; Thomas Hurka and Esther Shubert, "Permissions to Do Less Than the Best: A Moving Band," in Mark Timmons, ed., *Oxford Studies in Normative Ethics, Volume 2* (Oxford: Oxford University Press, 2012), 1–27.

Fiona Woollard have shown.²³ They have both advanced arguments for constraints against doing harm that are structurally very similar to my argument here—they focus on one’s ownership of one’s body, where I focus on one’s authority over one’s well-being. We can adapt their arguments to generate constraints against harming others that have just the same foundations as our agent-centered options. That people have special authority over how their well-being features in the justification of their actions and the actions of others not only grounds agent-centered options, but also gives them rights not to be harmed for the greater good. Indeed, I think these points are inextricable from one another: if they lacked the option not to sacrifice themselves, then in most cases it would be permissible to impose these costs on them. The argument from self-ownership therefore has adequate internal resources to rebut Kagan’s worry that we will end up justifying self-favoring options to harm others.

With self-sacrificing options, matters are simpler still. Provided my self-sacrifice affects only my self-regarding interests, or only those interests and others’ “me-regarding” interests, then it is permissible. If it contravenes some other value, for example by displaying inadequate self-respect, or by realizing an objectionable kind of relationship, then one’s right to a say can be outweighed by these other considerations.

One could also question whether this account adequately accommodates agent-centered options to promote the interests of those with whom one shares a valuable relationship. Perhaps I may forego pursuing the greater good not merely because doing so would be costly for me, but because it would adversely affect someone I am closely connected to—my son, for example. This seems like a standard case of an agent-centered option, but it is unclear how my view would vindicate it. It is my *son’s* well-being that is at stake in this choice, not my own.

I think this worry is misplaced. A full examination of our reasons of partiality is beyond the scope of this essay, but here’s a rough sketch. I think that reasons of partiality can cover the full gamut: sometimes acting on them is both all things considered morally best, and morally required; sometimes it is morally best, but not morally required; and sometimes it is not morally best, but nonetheless is permissible. Each of these scenarios involves one being permitted not to do what would be best, if we set your partiality aside. But only the second and third involve genuine agent-centered options. And in each of those, the action is made optional by the severe cost involved to the agent’s self-regarding interests.

To complete the sketch, we need a simple picture of the kinds of reasons that make an action morally best, all things considered. In particular,

²³ Quinn, “Doing and Allowing”; Woollard, “If This Is My Body.”

we need to distinguish between agent-relative and agent-neutral reasons. Agent-relative reasons include some pronominal back-reference to the identity of the agent acting on them.²⁴ They either apply only to that agent, or apply to her with special force. Agent-neutral reasons apply to all agents, with equal force, regardless of their identity.

Our valuable relationships give us agent-relative (as well as agent-neutral reasons). Sometimes, in virtue of those agent-relative reasons, acting partially is all things considered morally best, despite not being agent-neutrally morally best. If the cost to the agent is one she can be required to bear, for the sake of those reasons, then she is morally required to act partially. That is, she is prohibited from choosing the agent-neutrally best action.

If the cost exceeds what she is required to bear, then she might have an option to choose the agent-neutrally best action, and acting partially might be supererogatory. Think, for example, of circumstances in which you have a chance of saving your child's life only if you knowingly sacrifice your own.

Sometimes, however, the partial action is not all things considered morally best, even taking agent-relative reasons into account. Nonetheless, the agent is not required to choose the best action, because the cost to her of failing to act partially is great enough that she cannot reasonably be expected to bear that sacrifice.

Suppose you grant, then, that the argument is valid, and that it justifies a plausible set of agent-centered options—or at least, it meaningfully contributes to a pluralist defense of agent-centered options. One might still revert to Kagan's worry that intrinsic arguments for agent-centered options just end up redescribing the belief that we have options, rather than genuinely arguing for it. I agree that, this close to moral bedrock, it is hard to dig much deeper. But the idea that grounds agent-centered options, on my account, is an independently plausible one. If some fact that makes the world a better place is ineluctably and exclusively part of your life, then you should have some special authority over how that fact features in the justification of your actions and those of others. That is not just a restatement of the belief that we have agent-centered options.

V. CONCLUSION

In this essay I have offered one argument for agent-centered options. It satisfies most of the desiderata set out in Section II: it does not over-generate options; it grounds genuine options, rather than requirements;

²⁴ David McNaughton and Piers Rawling, "Value and Agent-Relative Reasons," *Utilitas* 7, no. 1 (1995): 31–47; Philip Pettit, "Universalizability without Utilitarianism," *Mind* 96, no. 381 (1987): 74–82.

it grounds agent-centered options in something deeper—the fact that my well-being makes the world a better place only by affecting my life; and it justifies both self-favoring and self-sacrificing options. It arguably under-generates options, though it grounds some options that the positive argument of Williams and others fails to support (self-sacrificing options, and options where relatively trivial interests are at stake). Anyway, if we are pluralists about what justifies options, then it is unsurprising that no single account can, on its own, support the full range of options that commonsense morality would endorse. We need to consider them all in concert, when working out which costs an agent can reasonably be required to bear.

Philosophy, Australian National University