

Are Moral Reasons Morally Overriding?

Douglas W. Portmore

Abstract In this paper, I argue that those moral theorists who wish to accommodate agent-centered options and supererogatory acts must accept both that the reason an agent has to promote her own interests is a nonmoral reason and that this nonmoral reason can prevent the moral reason she has to sacrifice those interests for the sake of doing more to promote the interests of others from generating a moral requirement to do so. These theorists must, then, deny that moral reasons morally override nonmoral reasons, such that even the weakest moral reason trumps the strongest nonmoral reason in the determination of an act's moral status (e.g., morally permissible or impermissible). If this is right, then it seems that these theorists have their work cut out for them. It will not be enough for them to provide a criterion of rightness that accommodates agent-centered options and supererogatory acts, for, in doing so, they incur a debt. As I will show, in accommodating agent-centered options, they commit themselves to the view that moral reasons are not morally overriding, and so they owe us an account of how both moral reasons and nonmoral reasons come together to determine an act's moral status.

Keywords Agent-centered options · Imperfect reasons · Moral reasons · Morality · Nonmoral reasons · Overridingness · Rational options · Rationality · Supererogation

Many moral theorists hold both that agents have, in many instances, an agent-centered option—that is, the moral option of either promoting their own interests or sacrificing those interests for the sake of doing more to promote the interests of others—and that, in such instances, doing more to promote the interests of others is supererogatory. I will argue that these moral theorists must further accept both that the reason an agent has to promote her own interests is a nonmoral reason and that this nonmoral reason can prevent the moral reason she has to sacrifice those interests for the sake of doing more to promote the interests of others from generating a moral requirement to do so. These theorists must, then, deny that moral reasons morally override nonmoral reasons, such that even the weakest moral reason trumps the strongest nonmoral reason in the determination of an act's moral status (e.g., morally permissible or impermissible). If I am right, if these moral theorists are committed to the view that nonmoral reasons are relevant to determining whether or not an act is morally permissible, then it would seem that they have their work cut out for them. Not only will they need to determine what moral reasons there are, but also what nonmoral reasons there are and which of these are relevant to determining an act's moral status. Furthermore, they will need to account for how these two very different sorts of reasons—moral and nonmoral reasons—

D. W. Portmore

Department of Philosophy, Arizona State University, P.O. Box 874102, Tempe, AZ 85287

e-mail: douglas.portmore@asu.edu

“come together” to determine an act’s moral status.¹ I will not attempt to do this work here, but only to argue that the work needs to be done.²

I begin the paper by clarifying what it means to say that moral reasons are not morally overriding and by drawing a distinction between moral reasons and morally relevant reasons. Then, in Section 2, I argue that those who wish to accommodate agent-centered options must deny that moral reasons are morally overriding and claim that the nonmoral reason an agent has to promote her own interests can prevent the moral reason she has to sacrifice those interests for the sake of doing more to promote the interests of others from generating a moral requirement to do so. And, in Section 3, I argue that those who wish to accommodate many typical instances of supererogation must do the same. In Section 4, I rebut Shelly Kagan’s objection that if we take this approach to accommodating agent-centered options and supererogatory acts, we are forced to accept that agents are rationally required to promote their own interests and are, thus, rationally forbidden from performing the supererogatory alternative.³ I end, in Section 5, with a few concluding remarks concerning the implications of the view that moral reasons are not morally overriding.

1 Some Quick Clarifications

To say that one type of reason, say, *m*-reasons, overrides another, say, *n*-reasons, with respect to a certain kind of normative status, *N*, is to say that, in any situation where both types of reasons are present and an act, *x*, has a certain *N*-status, no modification of the situation that involves affecting only what *n*-reasons there are will change *x*’s *N*-status. That is, if *m*-reasons override *n*-reasons with respect to an act’s *N*-status, then even the weakest *m*-reason overrides the strongest *n*-reason in the determination of that act’s *N*-status. To illustrate, suppose that I am morally required to keep my promise to meet with some student, that the reason I have to keep such promises is a moral reason, and that the reason I have to further my own self-interest is a nonmoral reason. If moral reasons override nonmoral reasons with respect to an act’s moral status, then no modification of this situation where only the nonmoral reasons I have for doing something else are altered will change the fact that I am morally required to meet with the student. Thus even if the situation changes such that someone is now offering me a million dollars to give a lecture during the time at which I promised to meet with the student, I will still be morally required to meet with the student and forego this unique opportunity to make a million dollars. If moral reasons are morally overriding, then the nonmoral reason I have to make a

¹ For one possible account, see Chang (2004).

² I attempt to do some of this work in Portmore (2008b).

³ See Kagan (1994), which is in reply to Bratman (1994). See also Kagan (1991), which is in reply to Slote (1991).

million dollars is, no matter how strong, powerless to prevent the moral reason I have to keep my promise from generating a moral requirement to do so.

Note that any thesis to the effect that one type of reason overrides another must be indexed to a certain kind of normative status. And so we must distinguish the thesis that moral reasons are *morally* overriding from the thesis that moral reasons are *rationally* overriding. According to the latter, moral reasons always override nonmoral reasons in the determination of an act's rational status. If this thesis were true, then it would always be *objectively irrational* (i.e., contrary to reason) to refrain from doing what one has most moral reason to do, even if what one has most moral reason to do is only supported by the most trivial of moral reasons and opposed by the weightiest of nonmoral reasons.⁴ I find this thesis incredible, but I will not be arguing explicitly against it here. It is worth pointing out, though, that if moral reasons are not morally overriding, then this supports the contention that they are not rationally overriding either, for it would be quite odd if *moral* reasons were rationally, but not morally, overriding. In that case, we would always be rationally required, although not always morally required, to act in accordance with what we have most moral reason to do. Surely, though, if nonmoral reasons can prevent moral reasons from generating moral requirements, then they can also prevent moral reasons from generating rational requirements. In any case, the focus here will be on whether moral reasons are *morally* overriding.⁵

I will argue that any moral theorist committed to accommodating agent-centered options must deny that moral reasons are morally overriding. Thus, such theorists are committed to the view that nonmoral reasons can and do affect an act's moral status. But let me stave off one immediate objection. Some will say that, by definition, a moral reason is any reason that is relevant to determining an act's moral status, and thus it is conceptually impossible for a nonmoral reason (a reason that is not relevant to determining an act's moral status) to be relevant to determining an act's moral status. I suppose that one could define a 'moral reason' in this way, but I will adopt a different definition so as to preserve an

⁴ An act is *objectively irrational* if and only if the agent has decisive reasons not to perform that act. Thus, the objective rational status of an act is purely a function of the reasons for and against it and its alternatives, regardless of whether or not the agent is aware of them. By contrast, the subjective rational status of an act depends, not on what reasons there are, but on what reasons the agent takes there to be, or, alternatively, on the practical mental functioning of the agent—see Parfit (2008) and Gert (2004), respectively.

⁵ A separate question is whether an agent can be morally required to perform an act that she does not have most reason to perform, all things considered. Sarah Stroud argues that the answer is 'No', and she calls her thesis the "Overridingness Thesis"—see Stroud (1998). What Stroud calls the "Overridingness Thesis" is more often called *moral rationalism*, but, again, this thesis is distinct from the one that I'll be discussing: the thesis that moral reasons are morally overriding. There is, however, the following interesting relationship between the two. If moral rationalism is true, and if moral reasons are not rationally overriding, then moral reasons won't be morally overriding, for, in that case, what we can morally obligated to do will be limited to those acts that we have most reason to perform, all things considered.

important distinction between reasons that are relevant to determining an act's moral status and reasons that, morally speaking, count in favor of, or against, performing some action. I will call any reason that is relevant to determining an act's moral status a *morally relevant reason*, and I will call any reason that, morally speaking, counts in favor of, or against, performing some action a *moral reason*. This is an important distinction to make, because it may be that not all morally relevant reasons are moral reasons. It is possible that some reasons can justify performing acts that it would otherwise be morally impermissible to perform without themselves counting in favor of performing those acts, morally speaking. For instance, we might think that the fact I would personally gain from breaking a promise is not a moral reason to do so in that it does not, *morally speaking*, count in favor of my doing so.⁶ Nevertheless, the reason that I have to act for personal gain may be a morally relevant reason, for, perhaps, such a nonmoral reason could, if sufficiently weighty, justify my breaking a promise. In any case, it would be a mistake to rule out, by definition, the very real possibility that nonmoral reasons (i.e., reasons that, morally speaking, do not count in favor of, or against, any action) are relevant in determining an act's moral status.

Having clarified the relevant terminology, I now proceed to argue that those who wish to accommodate agent-centered options must deny that moral reasons are morally overriding.

2 Moral Reasons, Overridingness, and Agent-Centered Options

Consider the following typical instance of an agent-centered option. An agent has a certain sum of money that she can use either to secure a considerable benefit for herself or to secure a far more considerable net benefit for various needy, distant strangers. Suppose, for instance, that she must choose to use the money that she has saved either to place a down payment on a new home or to help various needy, distant strangers by donating it to Oxfam. In this and many other typical instances of agent-centered options, the following four claims hold:

- (1) The agent has the choice to act either self-interestedly or altruistically—that is, she has the choice either to promote her own self-interest or to sacrifice her self-interest for the sake of doing more to promote the interests of others.
- (2) It is morally permissible for her to act self-interestedly.
- (3) It is also morally permissible for her to act altruistically.
- (4) The reason that she has to act altruistically has sufficient *moral requiring strength* that it would, absent countervailing reasons, generate a moral requirement to act altruistically.

⁶ Nor does the fact that I would personally gain by performing the act count against my performing it, morally speaking.

Claim (4) is in need of both clarification and justification. I'll start by clarifying what the moral requiring strength of a reason is. Reasons seem to play at least two normative roles in determining an act's moral status: a requiring role and a justifying role. Reasons can morally require performing acts that it would otherwise be morally permissible to refrain from performing (the requiring role), and reasons can morally justify performing acts that it would otherwise be morally impermissible to perform (the justifying role). Given these two normative roles, reasons have two potentially separable dimensions of strength: *moral requiring strength* and *moral justifying strength*. A reason has moral requiring strength to the extent that it can make it morally impermissible to refrain from performing acts that it would otherwise be morally permissible to refrain from performing, and a reason has moral justifying strength to the extent that it can make it morally permissible to perform acts that it would otherwise be morally impermissible to perform. We can, then, determine whether one reason, *R1*, has more moral requiring strength and/or more moral justifying strength than another, *R2*, using the following criteria:⁷

R1 has more *moral requiring strength* than *R2* if and only if:

- (i) *R1* would make it morally impermissible to do anything that *R2* would make it morally impermissible to do.
- (ii) *R1* would make it morally impermissible do some things that *R2* would *not* make it morally impermissible to do.

R1 has more *moral justifying strength* than *R2* if and only if:

- (i) *R1* would make it morally permissible to do anything that *R2* would make it morally permissible to do.
- (ii) *R1* would make it morally permissible do some things that *R2* would *not* make it morally permissible to do.

To illustrate the former criterion, consider that, on commonsense morality, the reason one has to refrain from killing an innocent person has more moral requiring strength than the reason one has to prevent an innocent person from dying. This is true in virtue of the following two facts: (i) If it would be immoral to do something—e.g., to check one's email—because it would entail failing to prevent an innocent person's death, then it would also be immoral to do that same thing—i.e., to check one's email—if it would entail killing an innocent person. (ii) Even though it would be morally permissible to let an innocent person die in order to save one's daughter (as where both are drowning and one has only enough time to save one of the two), it would not be morally

⁷ These are adapted from Joshua Gert's criteria for *rational* requiring strength and *rational* justifying strength. See Gert (2003, pp. 15-16).

permissible to kill an innocent person in order to save one's daughter (as where one's daughter needs that person's heart to live).

To illustrate the latter criterion, consider that, on commonsense morality, the reason one has to save three lives has more moral justifying strength than the reason one has to save one life. This is true in virtue of the following two facts: (i) If it would be morally permissible to do something that would otherwise be immoral—e.g., to break a promise to meet with a student—in order to save a life, then it would also be permissible to do that same thing—i.e., to break that promise to meet with the student—in order to save three lives. (ii) Even though it would not be morally permissible to fail to save two lives in order to save just one (assuming that everything else is equal and, thus, that there is no reason that favors saving the one as opposed to the two), it would be morally permissible to fail to save two lives in order to save three.⁸

Having clarified what is meant by the phrase 'moral requiring strength' in (4), I now need to justify (4). To see that the reason the agent has to act altruistically does have sufficient moral requiring strength that it would, absent countervailing reasons, generate a moral requirement to act altruistically, consider the following two cases. In the first case, a woman named Fiona is accessing her savings account via the Internet, and is about to transfer the entire balance to her escrow company so as to place the necessary down payment on a new home. She must do this if she is to purchase a new home, and she can do this simply by clicking on button A. However, there is an alternative. By clicking on button B, her savings will be transferred, not to her escrow company, but to Oxfam. By clicking on button B, then, she will be providing various needy, distant strangers in the Third World with some considerable benefit (e.g., potable water). Those who accept that there is an agent-centered option in such cases believe that, given the tremendous sacrifice involved, Fiona is not morally required to click on button B.⁹ But they also believe that the fact that her doing so would produce such a considerable benefit for these distant, needy strangers constitutes a reason of considerable moral requiring strength to click on button B. Indeed, but for the costs involved, it seems that this reason would generate a moral requirement to click on button B.

To see this, consider the second case, a variant of the first. In this case, Fiona can transfer the money to her escrow company by clicking on *either* button A or button B, and, in this case, a very rich man has agreed to transfer an equivalent sum of his own money to Oxfam if, and only if, Fiona clicks on button B. So, in

⁸ Of course, John Taurek has argued that it might be permissible to save the one instead of the two if one were to flip a coin to decide which group to save—see Taurek (1977). Nevertheless, I'm claiming only that, on commonsense morality, it would be impermissible to save the one instead of the two absent some special reason for saving the one.

⁹ If the reader believes that, given that lives are at stake, Fiona is morally required to click on button B, then imagine a revised version of the case where the net benefit that the strangers would receive were Fiona to click on button B is only slightly more significant than the net benefit that Fiona would receive were she to click on button A.

either case, Fiona will purchase her new home, but, by clicking on button B, she will also secure a considerable benefit for various others. Assume that there are no other morally relevant facts. Surely, in this case, Fiona is morally required to click on button B, for there is no good reason why she should not do so. By clicking on button B, she can purchase her new home while also providing an even more considerable net benefit for a number of others, and she can do so at no cost to herself, at minimal cost to the rich man (who, given the diminishing marginal utility of money, has more money than he can effectively use to benefit himself), and at absolutely no cost to anyone else. If one thinks that beneficence is only required when the would-be beneficiary is in great need or below a certain threshold of well-being, assume that those various needy, distant strangers that will be helped by Fiona's clicking on button B are below this threshold.

Given that the advocate of agent-centered options thinks that the reason Fiona has to click on button B gives rise to a moral requirement in the absence of countervailing reasons, she must conclude that it is a reason of sufficient moral requiring strength to generate a moral requirement in the absence of countervailing reasons. The only relevant difference between this case—the case of Fiona's clicking on button B and helping the stranger at no cost to herself—and the first case is how costly it is for Fiona to provide the strangers with a considerable benefit. But surely it is implausible to suppose that Fiona's reason to benefit the strangers or its moral requiring strength diminishes as the cost of her doing so increases.¹⁰ Suppose, for instance, that we were to gradually increase the cost of clicking on button B, from no cost at all, to 10 cents, to 20 cents, to 30 cents, etc.¹¹ It is not as if there is less and less to be said in favor of providing the strangers with a considerable benefit as the cost of clicking on button B increases

¹⁰ See Kagan (1989, p. 49). Of course, particularists might object that a reason can have a great deal of moral requiring strength in one context (e.g., the context in which the cost of acting altruistically is quite low) but very little to no moral requiring strength in another context (e.g., the context in which the cost of acting altruistically is quite high), for particularists will deny that reasons have any stable valence or strength values across possible contexts. But see Gert (2007) for an interesting and powerful response. Gert argues that the particularist "cannot merely deny that it makes sense to ascribe *stable* strength values to reasons: values that they keep from context to context. Rather, the particularist must make the blanket claim that talk of the strength of a reason makes no real sense even in a restricted context" (2007, p. 553). This is because, when we assign a strength value to a reason, we are providing a concise representation of the way it affects the normative statuses of acts across a range of contexts. For instance, if we assign greater moral requiring strength to one reason than another, we are committed to a claim about how these two reasons affect the moral statuses of acts across contexts: specifically, we must claim that the one reason would make it morally impermissible to do anything that the other reason would make it morally impermissible to do. Since this is precisely the sort of claim that particularists must deny, they must deny that talk of the strength of a reason makes any sense at all, even in particular contexts.

¹¹ If this just seems like an instance of Sorites Paradox to you, then increase the size of the incremental changes to the point where vagueness is no longer an issue. I don't see how the size of the increases matters. It seems to me that even if we increase the cost to the agent in one thousand dollar increments, it still won't feel like there is less and less to be said in favor of her benefiting the strangers with each incremental increase. Thanks to Dale Dorsey for raising this worry.

for the agent. At least, that is not what the phenomenology of the case tells us, for it feels like a case where one reason is outweighed by another, not like a case where one reason is undermined by another. If it were the latter, then once the cost was high enough, Fiona should cease to feel any pull toward clicking on button B. But even when the cost is extremely high, it is still clear that the fact that clicking on button B would result in a considerable benefit for a number of others speaks in favor of Fiona's doing so, morally speaking. So we should conclude that those who are committed to typical instances of agent-centered options are committed not only to claims (1)-(3), but also to claim (4).

Given (4), we must ask: What prevents the moral reason the agent has to act altruistically, with its considerable moral requiring strength, from generating a moral requirement to act altruistically? Clearly, it must be the reason the agent has to act self-interestedly, as this is the only countervailing reason, and we must cite some countervailing reason, since, given (4), we are to assume that the moral reason the agent has to act altruistically would generate a moral requirement to do so absent countervailing reasons. We must also assume that this countervailing reason to act self-interestedly must have at least as much moral justifying strength as the reason the agent has to act altruistically has moral requiring strength; indeed, it must have significantly more, because we think that the option to act self-interestedly would remain even if there was an increase in the moral requiring strength of the reason that favors acting altruistically. Lastly, we must assume that this reason to act self-interestedly must have less moral requiring strength than moral justifying strength, for, otherwise, we would end up with a moral requirement to act self-interestedly instead of a moral option to act either altruistically or self-interestedly. This is Shelly Kagan's worry. He says,

If, in some particular case, the balance of morally relevant reasons did not favor promoting the overall good [i.e., acting altruistically] but favored instead promoting the agent's own interests [i.e., acting self-interestedly]—then it seems that these reasons would still go on to generate a moral *requirement*. Admittedly, the agent would not be morally required to promote the overall good, but she *would* be morally required to promote her interests. Yet...[w]hat we were looking for was a defense of a moral *option*, according to which the agent would still be morally permitted (although not required) to do the act with the best results overall. (1994, pp. 338-9)

The solution, as a number of philosophers have pointed out (Slote, 1991; Bratman, 1994), lies in claiming that the morally relevant reasons that favor acting self-interestedly as opposed to altruistically are nonmoral reasons. If such nonmoral reasons can prevent the moral reason one has to act altruistically from generating a moral requirement, then what we end up with is a moral option rather than a moral requirement to act self-interestedly, since nonmoral reasons, by definition, have no moral requiring strength.¹² Kagan overlooks this possible

¹² If a given reason for action did have some moral requiring strength, it would thereby count in favor of performing that act, morally speaking, and would, therefore, be a moral reason.

solution to his worry, because he assumes that moral reasons are morally overriding and that, therefore, the only sorts of reasons that could prevent the moral reason one has to act altruistically from generating a moral requirement to do so is a moral reason to do something else. He says, “since we are concerned with what is required by *morality*, the relevant reasons—whether decisive or not—must be moral ones” (1989, p, 66). But Kagan’s inference is unwarranted; we should not just assume that moral reasons are morally overriding such that nonmoral reasons are just irrelevant with regard to what is required by morality.

Fortunately for the defender of agent-centered options, it is quite plausible to suppose that the reason one has to further one’s self-interest is a nonmoral reason and is, thus, a reason with no moral requiring strength. Moral reasons are, of course, a proper subset of reasons for action. So if reasons are considerations that count in favor of, or against, performing an action, then moral reasons are considerations that, *morally speaking*, count in favor of, or against, performing an action. But there is nothing, *morally speaking*, that counts in favor of promoting one’s self-interest, as such.¹³ This is not to say that one never has a moral reason to do what will further one’s self-interest—one sometimes does, as when doing one’s moral duty coincides with promoting one’s self-interest. The claim, then, is only that the mere fact that performing some act will further one’s self-interest does not itself, *morally speaking*, count in favor of doing so, for the mere fact that performing some act would be in one’s self-interest is never enough to make an act morally obligatory, or even morally supererogatory. The fact that I would benefit from getting a massage does not, *morally speaking*, count in favor of my getting one. If I had the opportunity to get one for free and chose instead to do something less beneficial for myself, I could rightly be called foolish or imprudent, but not immoral. Thus the reason one has to act self-interestedly doesn’t seem capable of making an act morally obligatory. Nor is it capable of making an act morally supererogatory. Consider, for instance, that in those instances where I am morally required to come to someone’s aid, as where there is a child drowning in a shallow pond, ensuring that I benefit myself in the process (by, say, alerting the news media so that I might receive some reward) would not count as an instance of going above and beyond the call of duty.

Of course, someone might object that there are duties to the self, and that such duties show that there is a moral reason to promote one’s self-interest. But the idea that there are certain duties to the self is compatible with the claim that there is no moral reason, *per se*, to promote one’s self-interest. To illustrate, take the duty to develop one’s talents. It seems that this duty derives, not from a duty to pursue prudential goods (i.e., one’s self-interest), but from either a duty to pursue certain perfectionist goods or a duty not to waste valuable gifts, for we are not morally obligated to develop every talent that would be of prudential

¹³ For a more thorough defense of this claim than what appears below, see Portmore (2003, Section III).

benefit to ourselves.¹⁴ Take, for instance, the ability to walk on one's hands over great distances. This is not the sort of talent that one is morally obligated to develop. Of course, one might benefit from developing such a talent, as where one wishes to make it into *The Guinness Book of World Records*. But even then, one isn't morally required, but only prudentially required, to develop that talent. This suggests that the reason it is wrong to waste certain talents has nothing to do with the individual's self-interest. If that were the case, we should object to the failure to develop any talent that would be of potential benefit to the individual who has it. But we don't; we don't morally object to wasting one's talent for walking on one's hands even where developing that talent would promote one's self-interest. Consider also that it seems that it would have been wrong for Mozart to have wasted his unique musical gifts even if he would have been slightly better off (prudentially speaking) doing something else. It seems, then, that the wrongness of wasting such great gifts lies with its wastefulness and not with its effects on the individual's self-interest. So we can admit that people are sometimes required to develop their talents, but we shouldn't infer from this that there is a moral reason to promote one's self-interest.

Of course, there is no denying that, on some moral theories (e.g., act utilitarianism), there is as much a moral reason to promote one's own interests as to promote anyone else's interests. Nevertheless, the point of this paper is to show that those moral theorists who endorse agent-centered options are committed to the view that the reason one has to promote one's self-interest is a nonmoral reason and that such a nonmoral reason can prevent a moral reason from generating a moral requirement. Since utilitarians deny the existence of agent-centered options, what they think just isn't relevant to the issue at hand. Also, unlike utilitarians, moral theorists who endorse agent-centered options typically give some credence to the sorts of intuitions that I've appealed to above. So my appeal to intuitions should hold some sway over them. In any case, though, it seems that such moral theorists must hold that the reason we have to promote our self-interest is a nonmoral reason if they wish to avoid, as Kagan worries, trading one moral requirement for another. And so we should conclude that those committed to agent-centered options are further committed to:

(5) The reason the agent has to act self-interestedly is a nonmoral reason.

But how can all of (1)-(5) be true? The only way is if the following is also true:

¹⁴ Perfectionist goods are not equivalent to, and do not necessarily correlate with, prudential goods. See Sumner (1996, pp. 23-4).

- (6) Moral reasons are not morally overriding—nonmoral reasons can, and sometimes do, prevent moral reasons, even those with considerable moral requiring strength, from generating moral requirements.¹⁵

Unless (6) is true, there is no way that (1)-(5) could all be true. To see this, consider the following indirect proof. If, contrary to (6), moral reasons are morally overriding, then the reason the agent has to act self-interestedly would be powerless to prevent the moral reason the agent has to act altruistically from generating a moral requirement to act altruistically, for, according to (5), the reason she has to act self-interestedly is a nonmoral reason. Clearly, if moral reasons are morally overriding, then nonmoral reasons, such as this one, would be powerless to prevent them from generating moral requirements. And, given (4), we must assume that the agent has a moral reason to act altruistically and that it has considerable moral requiring strength, such that it will generate a moral requirement absent countervailing reasons. Now, the only countervailing reason in this instance is the reason the agent has to act self-interestedly, but, as we've just established, this nonmoral reason is incapable of preventing the moral reason she has to act altruistically from generating a moral requirement. Thus, if we deny (6), we are forced to accept that the agent is morally required to act altruistically, and that would mean that we would have to deny (2)—that is, we would have to deny that it is morally permissible for the agent to act self-interestedly. So in order to accept all of (1)-(5), we must accept (6). And since those who accept typical instances of agent-centered options are committed to (1)-(5), they are also committed to (6). That is, they are committed to the view that moral reasons are not morally overriding.

In the next section, I approach the same problem from a different angle, showing that those who wish to accommodate typical instances of supererogation must accept that moral reasons are not morally overriding.

3 Moral Reasons, Overridingness, and Supererogation

Let me begin by stating two necessary conditions for an act's being supererogatory:

¹⁵ One way this might be true is if moral rationalism is true. Moral rationalism is the view that an agent can be morally required to perform a given act only if there is most reason, all things considered, to perform that act. If this is right, then a nonmoral reason to do something other than x could prevent a moral reason to do x from generating a moral requirement to do x by tipping the balance of reasons, all things considered, in favor of doing something other than x . Alternatively, one might suppose that the nonmoral reasons the agent has to do something other than x need not outweigh what moral reasons she has to do x in order to prevent it from generating a moral requirement to do x . Perhaps, these nonmoral reasons need only be sufficiently weighty, even if not weighty enough to outweigh the opposing moral reasons, to prevent these moral reasons from generating a moral requirement. I thank an anonymous reviewer from this journal for pointing out this other possibility to me.

S's performing *x* is supererogatory only if there exists some available alternative, *y*, such that:

- (a) *S* is morally permitted both to perform *x* and to perform *y*, and
- (b) *S* has more moral reason to perform *x* than to perform *y*.¹⁶

Although there is wide agreement among philosophers that a supererogatory act has to be, in some sense, morally superior to its non-supererogatory alternatives, some might deny that (b) is necessary, claiming instead that a supererogatory act need only be more morally praiseworthy than its non-supererogatory alternatives, and, on some accounts of praiseworthiness, an act can be more morally praiseworthy than another without there being more moral reason to perform it.¹⁷ While I do not want to deny (or assert) that a supererogatory act must be more morally praiseworthy than its non-supererogatory alternatives, I do want to argue that (b) is necessary.

To see why (b) is necessary, consider the following case. Suppose that Jane must choose either to buy an expensive pair of shoes to wear around town or to donate the money that she would have spent on those shoes to Oxfam, and suppose that both options are morally permissible. Further suppose that Jane falsely believes that buying a new pair of shoes would be what is best for others, whereas donating the money to Oxfam would be what is best for her. Suppose, then, that she fails to appreciate the moral reasons there are for her to donate the money to Oxfam, but that she thinks that she has a compelling self-interested reason to do so, for she believes that people will like her more (something she cares about) if she donates the money to Oxfam. She also falsely believes that she has good moral reason to buy the shoes, for she believes (falsely) that people will really enjoy seeing her wear them around town. So Jane decides, with altruism in her heart, to buy the shoes. Given her actual beliefs, motives, and intentions, one might claim that her buying the shoes is more morally praiseworthy than her donating the money to Oxfam. Yet, given that there is in fact no moral reason for her to buy the shoes and considerable moral reason for her to donate the money to Oxfam, it is hard to see how we could rightly call her act of buying the shoes

¹⁶ I leave open the question of whether there are any further necessary conditions, such as (c) *S*'s performing *x* is more morally praiseworthy than *S*'s performing *y*.

¹⁷ Someone else might deny (b), suggesting that a supererogatory act is one that involves a greater self-sacrifice for the sake of others than is required, whether or not there is necessarily any moral reason for agents to make such self-sacrifices—see, for instance, Hardwood (1998) and Vessel (2008). One problem with such an account is that it rules out the possibility of supererogation with respect to self-regarding duties. Yet it certainly seems possible to go above and beyond what such duties require—see Kawall (2003). Another problem is that if we deny that there is any better moral reason to go beyond what duty requires in terms of making self-sacrifices for others than to go beyond what duty requires in terms of, say, perspiration, then it's hard to see why only the former and not the latter would count as supererogatory. For more on this, see Portmore (2007, pp. 39-73).

an instance of going above and beyond the call of duty. Of course, the easy reply here is to insist that an act cannot be morally praiseworthy if the agent's choice of action is the result of her failure to appreciate the force of the moral reasons that favor it over its non-supererogatory alternatives, and I am quite sympathetic to this reply. But if one concedes that in order for an act to be morally praiseworthy the agent must properly appreciate the relevant moral reasons, then one must also concede that in order for an act to be more morally praiseworthy than another there will have to be more moral reason to perform it. And this is just to concede my point: that (b) is a necessary condition for an act's being supererogatory.

A further reason to accept (b) is that it nicely accounts for the normative force that supererogatory acts have, for there is a sense in which supererogatory acts are acts that agents morally ought to perform. It would be better, morally speaking, were they to perform them. Indeed, the facts that make an act supererogatory are presumably considerations that, morally speaking, count in favor of performing that act over any of its non-supererogatory alternatives. But if, contrary to (b), agents do not have more moral reason to perform a supererogatory act than to perform any of its non-supererogatory alternatives, then it is hard to see why it would be morally advisable for them to do so and why there is supposedly something that, morally speaking, counts in favor of their doing so.

Once we accept these two necessary conditions, however, supererogation can seem almost paradoxical since (b) appears to be in tension with (a), as James Dreier has explained:

Morality, we are inclined to think, is a matter of what reasons one has *from the moral point of view*. When there is a supererogatory act available, it would be better for you to perform it. So surely you have a reason, from the moral point of view, to perform the act. You may have some reason not to perform it, but at least typically you have no reason *from the moral point of view* [that is, no moral reason] to refrain from it (if you do have such reason, then it will ordinarily be outweighed by the reason you have to perform, because by hypothesis it is better to perform). But now it is hard to see how it could be permissible, from the moral point of view, to refrain from doing something that you have an undefeated reason (from that very point of view) to do. Everything from the moral point of view speaks in favor of your...[performing the supererogatory act], and nothing at all speaks against it. [In] what sense is it "all right," "permissible," "not wrong" to fail [to do so]? There seems to be no sense at all. Supererogation, according to this way of seeing things, turns out to be impossible. (2004, p. 148)

To dissolve the apparent paradox, the supererogationist must explain why the morally undefeated reason that favors performing the supererogatory act fails to generate a moral requirement to so act. There are only two possible explanations, for either moral reasons are morally overriding or they are not. If, on the one hand, they are, then the only possible explanation for why the morally undefeated reason that favors performing the supererogatory act fails to generate a moral requirement is that it has insufficient moral requiring strength to

generate a moral requirement.¹⁸ Call this *the insufficient-moral-requiring-strength explanation*. If, on the other hand, moral reasons are not morally overriding, then another possible explanation for why the morally undefeated reason that favors performing the supererogatory act fails to generate a moral requirement is that there is a nonmoral reason to perform some non-supererogatory alternative that prevents it from doing so. Call this *the nonmoral-reason explanation*. On the insufficient-moral-requiring-strength explanation, the assumption is that moral reasons are morally overriding, and so the thought is that the only possible explanation for why the moral reason that favors performing the supererogatory alternative fails to generate a moral requirement even when morally undefeated is that this moral reason lacks sufficient moral requiring strength to generate a moral requirement even when morally undefeated. On the nonmoral-reason explanation, by contrast, the thought is that moral reasons (even ones with considerable moral requiring strength) are not morally overriding, and thus the nonmoral reason the agent has to perform some non-supererogatory alternative can prevent the moral reason she has to perform the supererogatory alternative from generating a moral requirement to do so. On this view, then, what explains the fact that it is morally permissible for the agent to fail to perform the supererogatory alternative is the fact that she has a sufficiently weighty nonmoral reason to perform some non-supererogatory alternative.

Philosophers such as James Dreier and Michael J. Zimmerman overlook the nonmoral-reason explanation, because they assume (explicitly in Dreier's case and implicitly in Zimmerman's case) that moral reasons are morally overriding.¹⁹ On their view, it can never be morally permissible to do something that is supported by reasons of less moral requiring strength than those supporting some alternative, even if the agent has an undefeated nonmoral reason to do so and thus most reason, all things considered, to do so. Hence, they both take the

¹⁸ One might rightly point out that even a morally undefeated reason of considerable moral requiring strength will fail to generate a moral requirement when it is opposed by some moral reason of equal or incommensurate moral requiring strength. But this cannot explain why the morally undefeated reason that favors performing the supererogatory act fails to generate a moral requirement in those instances where it defeats whatever moral reasons there are for performing some permissible non-supererogatory alternative. Moreover, the morally undefeated reason that favors performing the supererogatory act must defeat (not just equal) these moral reasons for performing the permissible non-supererogatory alternative if it is to meet condition (b). Thus, assuming that moral reasons are morally overriding, the explanation for why the morally undefeated reason for performing the supererogatory alternative fails to generate a moral requirement has to be that it has insufficient moral requiring strength.

¹⁹ That Dreier thinks that an act's moral status is a function of solely moral reasons is clear from the first sentence in the above quote as well as from what he says on p. 149 of the same article. Zimmerman, by contrast, is less explicit, but he does say that if there being more moral reason to perform the supererogatory alternative is essential to supererogation, then any theory wishing to accommodate supererogation will have to declare that there are two sets of moral reasons, deontic and non-deontic reasons (or what I am calling moral reasons with, and moral reasons without, sufficient moral requiring strength)—see Zimmerman (1993, pp. 375-6).

insufficient-moral-requiring-strength explanation to be the only possible explanation.

In an effort to spell out how the insufficient-moral-requiring-strength explanation might go exactly, Dreier speculates that there might be two moral points of view, one the point of view of justice and the other the point of view of beneficence. Dreier further speculates that reasons stemming from justice have considerable moral requiring strength, but that reasons stemming from beneficence have no moral requiring strength. According to Dreier, supererogatory acts are more beneficent, but not more just, than their non-supererogatory alternatives.²⁰ So although agents have more moral reason to perform a supererogatory act than to perform any of its non-supererogatory alternatives, they are not morally required to do so, for the relevant reasons have no moral requiring strength (i.e., reasons of beneficence). What an agent is morally required to do is a function of only those reasons that have moral requiring strength (i.e., reasons of justice), and there is simply no more reason of this sort to perform the supererogatory alternative.

The problem with Dreier's proposed account and with the insufficient-moral-requiring-strength explanation in general is that it rests on the mistaken assumption that the reasons that make a supererogatory alternative morally superior to its non-supererogatory alternatives are always moral reasons of insufficient moral requiring strength. To the contrary, it seems that in many typical instances of supererogation the moral reason that favors performing the supererogatory alternative over its non-supererogatory alternatives is of considerable moral requiring strength. To illustrate, recall the case from section 2 where Fiona must choose between acting so as to secure a considerable benefit for herself by transferring the money from her savings account to her escrow company and acting so as to secure a more considerable benefit for various needy, distant strangers by instead transferring those funds to Oxfam. In this case, her forfeiting the chance to buy a new home and instead donating her savings to Oxfam is supererogatory. Now given that Dreier and Zimmerman insist on the insufficient-moral-requiring-strength explanation, they must presume that, in this case, the moral reason that favors Fiona's transferring the money to Oxfam as opposed to her escrow company is a moral reason of no moral requiring strength (i.e., a reason of beneficence). If the reason were instead a moral reason of sufficient moral requiring strength, then her donating the money to Oxfam would, on their view, be obligatory, not supererogatory, because, on their view, whatever nonmoral reason she has to purchase a new home is powerless to prevent the moral reason she has to donate the money to Oxfam from giving rise to a moral requirement to do so. But, as we have already seen in Section 2, the moral reason that favors her donating the money to Oxfam

²⁰ Similarly, in note 11, Zimmerman (1993) offers an example where he supposes that reasons of fidelity have considerable moral requiring strength but that reasons of beneficence have little to no moral requiring strength.

is a moral reason of considerable moral requiring strength; absent countervailing reasons, Fiona is morally *required* to donate the money to Oxfam. And, in general, it is just not plausible to suppose that reasons of beneficence have no moral requiring strength, for there are many other cases like this where the moral reason one has to perform some beneficent and supererogatory act is a moral reason of considerable moral requiring strength. So if we are to account for many typical instances of supererogation, we are going to have to accept the nonmoral-reason explanation, thereby accepting, contrary to Dreier and Zimmerman, that nonmoral reasons, and not just moral reasons, are relevant to the determination of an act's moral status.

4 Kagan's Objection

In the preceding two sections, I argued that those who want to account both for agent-centered options and for many typical instances of supererogation must accept that the reason an agent has to act self-interestedly is a nonmoral reason and that this nonmoral reason can prevent the moral reason she has to act altruistically from generating a moral requirement to do so. In response to philosophers who have taken this approach to defending agent-centered options (e.g., Michael Slote and Michael Bratman), Kagan raises the following objection:

Slote is arguing, in effect, that whenever the agent has a moral option, then from the rational point of view the reasons the agent has for favoring her own interests outweigh the reasons that support promoting the greater good.... But if this is so, then what if anything prevents these reasons from grounding a rational *requirement* to favor her interests in each such case? (1991, p. 927)

If the answer is that nothing would, then this seems to be an unacceptable result, for it would mean that in such cases "it would be rationally forbidden—irrational—to choose to do the morally preferable act" (1991, pp. 927-8). That is, it would be contrary to reason to perform the self-sacrificing, supererogatory alternative. And Kagan finds it "mysterious why it should be considered morally meritorious for an agent to sacrifice his interests for the greater good given that, on this approach, the balance of reasons actually opposes making such a sacrifice" (1989, pp. 378-9). So the worry is that, in taking this approach to defending agent-centered options, we would just be trading a moral requirement for a rational requirement and that this is unacceptable, in part, because it implies that performing the supererogatory alternative is rationally forbidden.²¹

Kagan's objection, as stated above, is a bit too quick, for, contrary to what Kagan says, Slote needn't claim that, from the rational point of view, the nonmoral reason the agent has to act self-interestedly *outweighs* the moral reason

²¹ The account of supererogation that I gave in Portmore (2003) did have the implication that all supererogatory acts are objectively irrational, and Michael Byron and Betsy Postow rightfully objected to it for this reason. See Byron (2005) and Postow (2005).

the agent has to act altruistically. Slote could instead claim that these two reasons are exactly equal in strength. If so, Slote could plausibly claim that the fact that there is just as much overall reason to act self-interestedly as to act altruistically accounts for the lack of a moral requirement to act altruistically. Moreover, since there is just as much overall reason to perform one alternative as the other, there will not only be a moral option, but also a rational option, to perform either. As convenient as this might be, the thought that we could account for all agent-centered options by claiming that in each case the relevant reasons are exactly equal in strength is quite implausible. As Kagan himself points out, this couldn't account for the wide range of options that we take there to be, for such ties would be extremely rare. To illustrate the problem, suppose that I could save some stranger from having to endure great physical pain by sacrificing \$500.²² If Slote were to suppose that the reasons for and against sacrificing the \$500 exactly balance out and that this is what accounts for both the moral and rational option to do either, then, if the tie were broken, there would cease to be an option. Suppose, for instance, the situation were to change slightly such that there was now slightly more reason for or against making the sacrifice, as where either the amount of pain that the stranger faces or the amount that I would need to sacrifice to prevent it increased slightly. In that case, it seems that we must admit that there is now most overall reason to perform one of the two alternatives. Yet we think that there are a wide range of rational options, such that making such an altruistic self-sacrifice remains rationally optional even if the reasons for or against making that self-sacrifice increase slightly. If there's an option both before and after the strength or number of reasons in support of one of the two alternatives increases, we cannot hope to account for such options by supposing that the relevant reasons exactly balance out.

So Kagan's objection amounts to the following. If the reasons in favor of each alternative are not exactly balanced out, then it seems that one or the other alternative will be the one that there is most overall reason to perform, and, if so, that alternative will be rationally required even if not morally required. Thus this strategy for defending agent-centered options involves trading a moral requirement for a rational requirement, when what we were looking for was both a moral, and a rational, option to act either altruistically or self-interestedly.

It's important to realize, though, that the problem of accounting for rational options given that it seems implausible to suppose that the relevant reasons always balance out is a difficult philosophical problem that faces anyone interested in accounting for what Joseph Raz (1999, p. 100) calls:

The Basic Belief: In most typical choice situations, the relevant reasons do not require the performance of one particular act alternative, but instead permit performing any of numerous act alternatives.

²² I borrow this example from Kagan (1989, pp. 374-5).

It seems, for instance, that I could now, in accordance with reason, do any of the following: continue to work on this paper, read the newspaper, prepare for my next lecture, play with my daughter, relax by watching TV, or volunteer for Oxfam. But how are we to account for this? On the face of it, it seems that there could be such rational options only if there were exactly equal reason to perform each optional act alternative, and yet it is hard to believe that such is the case. For instance, it is difficult to believe that I have just as much reason to watch TV as to either volunteer for Oxfam or work on this paper. After all, volunteering for Oxfam seems vastly superior to watching TV in terms of the amount of impersonal good that it would do, and working on this paper seems vastly superior to watching TV in terms of the amount of personal good it would do me. Moreover, the fact that, in many of these choice situations, the relevant act alternatives remain rationally optional even when there is an increase in the number and/or strength of the reasons in support of one of the alternatives shows that their optional status cannot be due, in the first place, to a perfect balance of reasons in support of each act alternative (Gert 2008, p. 14). Consider, for instance, that it would still be rationally permissible for me to continue working on this paper even if there was a slight increase in the strength or number of reasons that favored one of my other options, as where, say, Oxfam institutes a new policy of providing volunteers with a free, delicious lunch. So the puzzle is to explain how, in most choice situations, there could be so many optional act alternatives if, as seems to be the case, there is not exactly equal reason to do each of these alternatives.

It seems, then, that the problem of accounting for rational options is not any more serious for the defender of agent-centered options than it is for anyone else who accepts the basic belief. And if the problem is not specific to this defense of agent-centered options, then it is no objection to this defense that it encounters the same problem that anyone wishing to account for rational options faces. Moreover, since this problem is quite general, Slote, Bratman, and others can appeal to the same sorts of solutions that others appeal to in accounting for the basic belief. Now I don't have space here to argue for any particular solution to this problem, but I do want to point out that there are a number of possible solutions and that most, if not all, are compatible with this sort of defense of agent-centered options.

There are, at least, five possible solutions to the problem of accounting for rational options. First, we could adopt a satisficing conception of rationality, where it is rationally permissible to perform a given act provided the reasons supporting it are sufficiently weighty, irrespective of whether or not they are outweighed by countervailing reasons. Given such a satisficing conception of rationality, we could account for the basic belief by claiming that, in most choice situations, the reasons in support of numerous act alternatives are all sufficiently weighty. Thus the defender of agent-centered options could even admit that the

nonmoral reason that the agent has to act self-interestedly *outweighs* the moral reason the agent has to act altruistically (thereby preventing it from generating a moral requirement) and yet deny that it generates a rational requirement, for she can claim that the moral reason the agent has to act altruistically is sufficiently weighty so as to make acting on it rationally permissible even though it is outweighed by the nonmoral reason the agent has to act self-interestedly.²³

Second, we could hold that, in addition to the relations ‘better than’, ‘worse than’, and ‘equally good as’, there is also the relation ‘on a par with’.²⁴ If we were, then, to claim that, in most choice situations, the reasons in support of numerous act alternatives are all on a par, we could account for there being numerous optional act alternatives that remain optional even when the strength or number of the reasons in support of one of them slightly increases. For instance, if, in the case of agent-centered options, the nonmoral reason the agent has to act self-interestedly was, rationally speaking, on a par with the moral reason the agent has to act altruistically, we could account for a rational option to act either way and maintain that these two alternatives would remain optional even if the strength and/or number of reasons in support of just one of them were to increase slightly.²⁵

Third, we could claim that reasons for action have not only two dimensions of moral strength but also two dimensions of rational strength: rational requiring strength and rational justifying strength. Roughly speaking, a reason has rational justifying strength to the extent that it can make it rationally permissible to perform acts that it would otherwise be irrational to perform, and a reason has rational requiring strength to the extent that it can make it irrational to refrain from performing acts that it would otherwise be rationally permissible to refrain from performing (Gert 2003). If we were, then, to claim that, in most choice situations, some of the relevant reasons have a great deal more rational justifying strength than opposing reasons have rational requiring strength, we could thereby account for the basic belief. Interestingly, Joshua Gert argues that, in

²³ And, as mentioned in note 18, one could claim that the nonmoral reason the agent has to act self-interestedly need only be sufficiently weighty, even if not weighty enough to outweigh the opposing moral reasons, to prevent these moral reasons from generating a moral requirement. Either way, there will, on the satisficing view, be a rational option to act either altruistically or self-interestedly so long as the relevant reasons in favor of each alternative are sufficiently weighty.

²⁴ As Ruth Chang notes, many philosophers think that “if two items A and B are evaluatively comparable, then A must be better or worse than B, or A and B must be equally good. Call this the ‘Trichotomy Thesis’”—see Chang (2002). Chang rejects the Trichotomy Thesis and argues that, in addition to these three, A and B might be on a par. The difference between being on a par and being equally good is that A and B can be on a par and a small improvement in either A or B would not necessarily make the improved item better than the unimproved item—the improved item might still just be on a par with the unimproved item. By contrast, if A and B are equally good, then a small improvement in one or the other would necessarily make the improved item better than the unimproved item.

²⁵ Of course, we must deny that, morally speaking, they are on a par if we hope to account for the fact that acting altruistically is supererogatory.

comparison to self-interested reasons (or, at least, in comparison to those self-interested reasons that pertain to avoiding non-trivial harms), altruistic reasons have no rational requiring strength, but just as much rational justifying strength when the amount of harm at stake for others is comparable to the amount of harm at stake for the agent (2004, 141). Thus, on Gert's view, one is, for instance, rationally permitted to sacrifice one's own life to save another's even though one is not rationally required to sacrifice anything to save someone else's life. If this is right, then the defender of agent-centered options can claim that, given its considerable moral justifying strength, the nonmoral reason the agent has to safeguard her own self-interest prevents the moral reason she has to act altruistically from generating a moral requirement to act altruistically and does so without thereby generating a rational requirement to act self-interestedly, for this moral reason has, on Gert's view, insufficient rational requiring strength to make acting altruistically rationally required, but sufficient rational justifying strength to make acting altruistically rationally permissible.

Fourth, we could claim that, in most choice situations, the relevant reasons are either not at all, or only very roughly, comparable.²⁶ When two competing reasons are incomparable, they neither defeat one another nor exactly balance out (and, furthermore, they would not be on a par). And since whenever the competing reasons for various act alternatives fail to defeat each other, it accords with reason to perform any one of them, claiming that, in most choice situations, the relevant reasons are incomparable and thus fail to defeat each other, allows us to account for the basic belief. Moreover, the defender of agent-centered options can appeal to the following facts: the sorts of reasons that support supererogatory acts are typically altruistic (or, alternatively, impartial) reasons, the sorts of reasons that support their permissible non-supererogatory alternatives are typically self-interested (or partial) reasons, and it is not, on the face of it, implausible to suppose that these two types of reasons are not at all, or only very roughly, comparable. Indeed, Henry Sidgwick (1907), for instance, held that impartial reasons and self-interested reasons are wholly incomparable, and this led him to accept a kind of *dualism of practical reason*, where it is always objectively rational to do either what would be impartially best or what would be self-interestedly best. Recently, Derek Parfit (2008) has argued that some form of dualism about practical reason is correct, but that Sidgwick was wrong to think that impartial and self-interested reasons are wholly incomparable. Parfit argues

²⁶ If two types of reasons, x - R and y - R , are *wholly incomparable*, then there is, for *no pair* of token instances of x - R and y - R , some truth as to how the two compare—that is, as to whether the one is stronger than, weaker than, or equally strong as (or on a par with) the other. If, by contrast, two types of reasons, x - R and y - R , are only *very roughly comparable*, then there is, for only a *few pairs* of token instances of x - R and y - R , some truth as to how the two compare. Just how rough the comparability of the two types of reasons is in proportion to how few pairs there are for which there is some truth as to how the two compare. The notions of 'wholly incomparable' and 'roughly comparable' come from Parfit (2008).

that they are instead only roughly comparable. In either case, there would, in most instances, be a rational option to act either altruistically or self-interestedly.

Now I'm a bit unsure as to whether this approach to accounting for rational options is ultimately tenable for the defender of agent-centered options. If, in a given case, the relevant moral and nonmoral reasons are incomparable, then how could the nonmoral reasons be sufficiently weighty *in comparison to* the moral reasons to prevent them from generating a moral requirement?²⁷ But, perhaps, no comparison of strength is needed. Perhaps, nonmoral reasons prevent moral reasons from generating moral requirements whenever these nonmoral reasons are, by some non-comparative measure, sufficiently weighty. Such sufficiently weighty nonmoral reasons would, then, prevent even the strongest of moral reasons from generating a moral requirement. This might have some interesting implications as to what sorts of agent-centered options agents will have. Thus whether the defender of agent-centered options can ultimately endorse such an approach will depend on whether or not she finds these implications acceptable.

Fifth, we could account for the basic belief by supposing that the relevant reasons are *imperfect reasons*—reasons that do not support performing any specific act alternative, but instead support performing any of a number of act alternatives that would each bring about the same valuable end. Of course, it may seem that acting self-interestedly and acting altruistically cannot be plausibly interpreted as two ways of pursuing the same valuable end or as two ways of acting on the same reason. But things are not as they seem. If we take note of the fact that our agency is extended over time and that, at a more general level of description, the agent would be choosing between two courses of action, one that includes acting self-interestedly now and acting altruistically later and another that includes acting altruistically now and acting self-interestedly later, it does seem plausible to construe these two courses of action as two ways of trying to achieve the same thing: specifically, a reasonably choice-worthy life—that is, a life containing a reasonable balance of egoism and altruism. If we think that the any reasonably choice-worthy life will contain some reasonable proportion of moral goods to prudential goods, then it does not matter whether or not one performs an altruistic act now or later so long as, either way, one is pursuing a course of action that will result in a reasonable balance being struck over time. Since both courses of action can lead to one's living a reasonably choice-worthy life, and since one of these two courses of action involves acting altruistically now and the other involves acting self-interestedly now, I have the rational option of acting either altruistically or self-interestedly now.²⁸

To sum up, I've argued that Kagan's objection poses no more serious a problem for the defender of agent-centered options than it does for anyone else concerned to account for a wide range of rational options. Moreover, I've shown

²⁷ I thank G. Shyam Nair for raising this concern.

²⁸ For both a more thorough explication and an extended defense of this possible solution, see Portmore (2008a).

that there are a number of possible solutions to this problem and that the defender of agent-centered options can avail herself of most, if not all, of them. I haven't tried to assess whether any of these proposed solutions are adequate, and, if some are, which of them is the most plausible. My goal here is only to establish that whether an adequate solution can be found or not, the defender of agent-centered options needn't fear Kagan's objection. If an adequate solution can be found, the defender of agent-centered options can most likely employ it to meet Kagan's objection. If an adequate solution can't be found, then the defender of agent-centered options will have to admit, as Kagan worries, that there will be a rational option to act either altruistically or self-interestedly only in those rare instances in which the reasons supporting each alternative exactly balance out. But, in that case, the objection cuts no ice, for, if no adequate solution can be found, then everyone (whether they endorse agent-centered options or not) will have to admit that there are rational options only when the relevant reasons exactly balance out.

5 Conclusion

I have argued that those who want to account both for agent-centered options and for many typical instances of supererogation must accept that the reason an agent has to act self-interestedly is a nonmoral reason and that this nonmoral reason can prevent the moral reason she has to act altruistically from generating a moral requirement to do so. I have shown, therefore, that such theorists must deny that moral reasons are morally overriding. If this is right, then it seems that these theorists have their work cut out for them. It will not be enough for them to provide a criterion of rightness that accommodates agent-centered options and supererogatory acts, for, in doing so, they incur a debt. As I've shown, in accommodating agent-centered options, they commit themselves to the view that moral reasons are not morally overriding, and so they owe us an account of how both moral reasons and nonmoral reasons come together to determine an act's moral status. This is no small task. First, it requires providing both a substantive account of what moral reasons there are and a substantive account of what nonmoral reasons there are. Second, it requires providing a substantive account of how moral and nonmoral reasons function together to determine an act's moral status. With regard to this second task, we will need answers to the following sorts of questions: (1) "Is it ever wrong to do what one has most reason to do, all things considered?"; (2) "Is it ever wrong to do what one has most moral reason to do?"; and (3) "Is it ever permissible to do something that one has neither most moral reason to do nor most reason to do, all things considered, and, if so, when?" I do not mean to suggest that philosophers have not been working on these issues; they have.²⁹ But what is not generally appreciated is the importance of resolving these issues if we are going to have any hope of

²⁹ See, for instance, Chang (2004), Stroud (1998), Portmore (2008b); and McLeod (2001).

defending a moral theory that accommodates agent-centered options and supererogatory acts.

Acknowledgement For helpful comments and discussions, I thank Richard Arneson, Noell Birondo, Dan Boisvert, Campbell Brown, Michael Byron, Dale Dorsey, Nir Eyal, Joshua Glasgow, G. Shyam Nair, Derek Parfit, Betsy Postow, Mark Schroeder, David Shoemaker, Mark van Roojen, Michael J. Zimmerman, numerous anonymous referees, students in my Spring 2006 seminar entitled “The Limits of Morality,” and the audience at my 2006 Pacific APA colloquium.

References

- Bratman, M. E. (1994). Kagan on ‘The appeal to cost’. *Ethics*, 104, 325-332
- Byron, M. (2005, April). Alternative consequentialisms. (Paper presented at Lingnan University, Hong Kong)
- Chang, R. (2002). The possibility of parity. *Ethics*, 112, 659-688
- Chang, R. (2004). Putting together morality and well-being. (In P. Baumann & M. Betzler (Eds.), *Practical conflicts: New philosophical essays* (pp. 118-158). Cambridge, Cambridge University Press)
- Dreier, J. (2004). Why ethical satisficing makes sense and rational satisficing doesn’t. (In M. Byron (Ed.), *Satisficing and maximizing* (pp. 131-154). Cambridge: Cambridge University Press)
- Gert, J. (2003). Requiring and justifying: Two dimensions of normative strength. *Erkenntnis*, 59, 5-36
- Gert, J. (2004). *Brute rationality: Normativity and human action*. (Cambridge: Cambridge University Press)
- Gert, J. (2007). Normative strength and the balance of reasons. *Philosophical Review*, 116, 533-562
- Gert, J. (2008). Michael Smith and the rationality of immoral action. *The Journal of Ethics*, 12, 1-23
- Hardwood, S. (1998). Eleven objections to utilitarianism. (In L. Pojman (Ed.), *Moral philosophy: A reader* (pp. 179-92). Indianapolis: Hackett Publishing)
- Kagan, S. (1989). *The limits of morality*. (Oxford: Clarendon Press)
- Kagan, S. (1991). Replies to my critics. *Philosophy and Phenomenological Research*, 54, 919-928
- Kagan, S. (1994). Defending options. *Ethics*, 104, 348-351
- Kawall, J. (2003). Self-regarding supererogatory actions. *Journal of Social Philosophy*, 34, 487-498
- McLeod, O. (2001). Just plain ‘ought’. *The Journal of Ethics*, 5, 269-291
- Parfit, D. (2008). Climbing the mountain. Retrieved January 29, 2008, from http://individual.utoronto.ca/stafforini/parfit/parfit_-_climbing_the_mountain.pdf
- Portmore, D. W. (2008a). Imperfect reasons. Retrieved January 24, 2008, from http://www.public.asu.edu/~dportmor/Imperfect_Reasons_and_Rational_Options.pdf
- Portmore, D. W. (2008b). Dual-ranking act-consequentialism. *Philosophical Studies* (in press). DOI 10.1007/s11098-006-9059-7
- Portmore, D. W. (2007). Consequentializing moral theories. *Pacific Philosophical Quarterly*, 88, 39-73
- Portmore, D. W. (2003). Position-relative consequentialism, agent-centered options, and supererogation. *Ethics*, 113, 303-332
- Postow, B. Supererogation again. *Journal of Value Inquiry*, 39, 245-253
- Raz, J. (1999). *Engaging reason*. (Oxford, Oxford University Press)
- Sidgwick, H. E. (1907) *The method of ethics*. (London: Macmillan and Company)
- Slote, M. (1991). Shelly Kagan’s *The Limits of Morality*. *Philosophy and Phenomenological Research*, 54, 915-917
- Sumner, L. W. (1996). *Welfare, happiness, and ethics*. (Oxford: Oxford University)

- Stroud, S. (1998). Moral overridingness and moral theory. *Pacific Philosophical Quarterly*, 79, 170-189
- Taurek, J. (1977). Should the numbers count? *Philosophy and Public Affairs*, 6, 293-316
- Vessel, J.-P. (2008, March). Supererogation for utilitarianism. (Paper presented at the Pacific Division Meeting of the *American Philosophical Association*, Pasadena, CA)
- Zimmerman, M. J. (1993). Supererogation and doing the best one can. *American Philosophical Quarterly*, 30, 373-80