

CONSEQUENTIALISM AND THE SELF-OTHER ASYMMETRY

I. Consequentialism, Commonsense Morality, and the Self-Other Asymmetry

Unlike traditional act consequentialism (TAC), commonsense morality holds that it is often permissible to act in ways that are detrimental to one's own self-interest even when doing so would fail to maximize the good. On commonsense morality, I can, for instance, forgo my own greater benefit for sake of securing someone else's lesser benefit. Moreover, I am, on commonsense morality, often permitted both to forgo some benefit and to incur some harm so long as this doesn't adversely affect anyone else's interests (or having other morally relevant repercussions), whereas, on TAC, it is impermissible both to forgo some benefit and to incur some harm unless doing so produces some compensating amount of good for others (or some compensating amount of impersonal good).

On commonsense morality, then, there are not only *agent-favoring options* (options either to favor one's own interests or to sacrifice one's own interests for the sake of doing more to promote the overall good), but also *agent-sacrificing options* (options either to maximize the good or to sacrifice one's own interests for the sake of doing more to promote the good of others). That is, on commonsense morality, there is not only the option to give one's own interests greater weight, but also the option to give one's own interests lesser weight, or even no weight at all.

So there is an asymmetry between what we are permitted to do to ourselves and what we are permitted to do to others. Although it is permissible to be indifferent toward one's own welfare, it is not permissible to be indifferent toward the welfare of others. Although it is permissible to kill oneself in order to prevent five other comparable killings, it is not permissible to kill someone else in order to prevent five other comparable killings. Even though it is morally permissible to gratuitously cause oneself pain, it is not morally permissible to gratuitously cause someone else pain.

Along with Slote, I think that there is here a morally relevant distinction between the self and others and not just a morally relevant distinction between having consent and not having consent. It would be wrong for me to sacrifice Jill's arm to save Jack's thumb, other things being equal, even if Jill consents to my doing so. Yet it would not be wrong (although, perhaps, stupid and foolish) for me sacrifice my arm to save Jack's thumb, other things being equal. If you think that there is a duty not do anything that causes yourself permanent disability, then just imagine a slightly different case, where the choice is between a severe headache and a mild headache as opposed to being between sacrificing an arm and sacrificing a thumb. And consider that if I'm harmed as a result of my own negligence, it would be false to say that I consented to being harmed. Yet, despite the lack of consent, it seems

that I did nothing wrong so long as my negligent act didn't put anyone else at risk. But it does seem that I've done wrong if I've harmed someone else as a result of my negligent act.

The asymmetry seems to be this: whereas there is a non-derivative moral reason to promote the welfare of others, there seems to be no (non-derivative) moral reason at all to promote one's own welfare.

This self-other asymmetry is reflected in our language. For instance, the word 'benevolent' applies only to what we do for others. The words 'murder' and 'unjust' also reflect this asymmetry in that it seems to be a conceptual mistake to assert the possibility of murdering oneself or of treating oneself unjustly.

II. Self/Other Utilitarianism

Self/Other Utilitarianism (SOU): S's performing x is morally permissible if and only if there is no available act alternative that would produce both (i) more utility for others (i.e., those other than S) than x would and (ii) at least as much overall utility as x would (Sider, "Asymmetry and Self-Sacrifice," p. 128). (Actually, on Sider's formulation, (ii) reads: "more overall utility as x would." I'll explain below my reasons for preferring that (ii) read: "at least as much overall utility as x would.")

Let $U_s(x)$ = the utility that accrues to S if S does x , $U_{-s}(x)$ = the utility that accrues to others if S does x , and $U(x)$ = the overall utility that is produced if S does x .

<i>act</i>	$U_s(x)$	$U_{-s}(x)$	$U(x)$	<i>moral status</i>
a ₁	+70	+10	+80	merely permissible
a ₂	+20	+5	+25	impermissible
a ₃	+10	+15	+25	supererogatory
a ₄	-10	+20	+10	supererogatory and morally best

SOU is quite an unorthodox version of act consequentialism in that it ranks outcomes according to two auxiliary rankings and then takes the moral status of an action to be a function of some principal ranking that is in turn a function of the two auxiliary rankings. SOU ranks outcomes both in terms of utility for others and in terms of utility overall and holds that an act is permissible if and only if there is no available alternative act whose outcome contains both (1) more utility for others and (2) at least as much utility overall. Let's call an act-consequentialist theory that ranks outcomes in terms of two auxiliary rankings "dual-ranking act-consequentialism" or "DRAC" for short.

SOU accommodates both supererogatory acts and agent-sacrificing options. However, it does not accommodate agent-favoring options. But a theory that

accommodates all three is easily formulated, employing this kind of dual-ranking structure that SOU has.

III. Schefflerian Utilitarianism

Schefflerian Utilitarianism (SU): S's performing x is morally permissible if and only if there is no available act alternative that would produce both (i) more utility for others (i.e., those other than S) than x would and (ii) at least as much egoistically-adjusted utility, where we include everyone's utility but adjust the overall total by giving S's utility ten times the weight of anyone else's.

Let $U_s(x)$ = the utility that accrues to S if S does x , $U_{-s}(x)$ = the utility that accrues to others if S does x , $U_{+s}(x) = U_{-s}(x) + [10 \times U_s(x)]$, and $U(x)$ = the overall utility that is produced if S does x .

act	$U(x)$	$U_s(x)$	$U_{-s}(x)$	$U_{+s}(x)$	moral status
a ₁	17	7	10	80	merely permissible
a ₂	7	2	5	25	impermissible
a ₃	16	1	15	25	supererogatory
a ₄	19	-1	20	10	supererogatory and morally best

SU accommodates supererogatory acts and both agent-favoring and agent-sacrificing options. The option to perform a₁ or a₄ is an agent-favoring option, for the benefit for S in S's performing a₁ as opposed to a₄ is only 8, whereas the benefit for others in S's performing a₄ as opposed to a₁ is 10. The option to perform a₁ or a₃ is an agent-sacrificing option, for the sacrifice for S in S's performing a₃ as opposed to a₁ is 6, whereas the benefit for others in S's performing a₃ as opposed to a₁ is only 5.

IV. The *ad hoc* Objection

A. The Objection

We can easily see how each of SOU's two auxiliary rankings are evaluative, for it's plausible to construe SOU as holding that there is a deep asymmetry between self and others such that although there is moral value in an agent's promoting the utility of others, there is only non-moral value in an agent's promoting her own utility. In that case, we should construe the first auxiliary ranking—the ranking in terms of total utility for others, i.e., $U_{-s}(x)$ —as a ranking of outcomes in terms of their moral value and construe SOU's second auxiliary ranking—the ranking in terms of overall utility, i.e., $U(x)$ —as a ranking of outcomes in terms of their all-things-considered value (weighing together both

moral and non-moral value).¹ We might ask, though, “Along what evaluative dimension does SOU’s principal ranking order outcomes?” Curiously, it seems that SOU’s principal ranking is *not* a ranking of outcomes along any evaluative dimension: not moral value, not non-moral value, and not all-things-considered value. Rather, SOU’s principal ranking is just a ranking that’s a function of two evaluative rankings.² Of course, it is not uncommon for consequentialists to adopt a more sophisticated ranking principle (as compared to, say, that of HAU) in the hopes of reconciling consequentialism with our commonsense moral convictions. But what typically makes this more than just an *ad hoc* move is that these consequentialists appeal to certain axiological intuitions that motivate the move to adopt the more sophisticated ranking principle. But what axiological intuitions could possibly motivate the move from HAU to SOU, the move from a single-ranking structure to a dual-ranking structure? Speaking of the principal rankings that SOU generates for various agents, Ted Sider admits that they do not “correspond to any independently important axiological facts.” He says, “All I claim for their importance is their role in determining moral normative status” (1993, p. 128). In that case, though, the move to a dual-ranking structure seems to be entirely *ad hoc*, motivated solely by a desire to accommodate the self-other asymmetry.

B. My Response

I concede that the move from a single-ranking structure to a dual-ranking structure isn’t motivated by any axiological intuitions. Nevertheless, I will argue that the move is not *ad hoc*, for it is, in fact, motivated by more than just a concern to accommodate the self-other asymmetry. To succeed, I’ll have to show both: (i) that there is no moral reason to promote one’s own utility *per se*³ and (ii) that moral reasons are not morally overriding, that non-moral reasons can and do sometimes prevent moral reasons from generating moral requirements. I’ve argued for both claims elsewhere, but I will briefly

¹ All-things-considered value isn’t a kind of value in the sense that moral value and non-moral value are two kinds of value. So I’ll be using the phrase ‘all-things-considered value’ as short for ‘value, all things considered (i.e., considering both moral and non-moral value)’.

² This accounts for why I call theories like SOU “*dual-ranking* act-consequentialism.” Although it’s true that SOU takes the deontic status of an action to be a function of a single ranking, act-consequentialism must hold that the deontic status of an action is a function of some *evaluative* ranking or, if not a function of some single evaluative ranking, then a function of some set of evaluative rankings—see Portmore (forthcoming). So SOU is a dual-ranking theory in that it makes the deontic status of an action a function of not just a single evaluative ranking but of dual evaluative rankings.

³ This is not to deny that agents are sometimes morally required to do what promotes their own utility. Nor is this to deny that there are duties to self. The claim here denies only that agents have such duties in virtue of the effects that abiding by them will have on their own utility. For more on this issue, see Portmore (2003).

summarize these arguments below.⁴ I'll then explain (iii) how establishing (i) and (ii) will motivate the move to a dual-ranking structure.

1. My Argument for (i)

According to (i), the fact that doing *x* would further *S*'s interests is a nonmoral reason for *S* to do *x* and is, thus, a reason with no moral requiring strength (a reason has moral requiring strength to the extent that it has the ability to make acts that it would otherwise be permissible to refrain from performing morally required). 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.⁵ 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

⁴ See Portmore (2003, 2008).

⁵ For a more thorough defense of this claim than what appears below, see my "Position-Relative Consequentialism, Agent-Centered Options, and Supererogation," *Ethics* 113 (2003): 303-332, especially section III.

self-interest), but from a duty to pursue certain perfectionist goods, for we are not morally obligated to develop every talent that would be of prudential benefit to ourselves.⁶ 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.

To sum up, the argument for (i) is as follows. There are a number of cases (e.g., the getting a massage case) where the mere fact that performing some act would promote the agent's self-interest does not count as a moral reason to perform that act. And there seem to be no cases where there is some moral reason to do what is one's self-interest but we can't account for this moral reason by appealing to some fact other than the mere fact that it is in one's self-interest. Thus, the best explanation for this data is (i).

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 here 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.

2. My Argument for (ii)

My argument for (ii) is the following: There are agent-centered options such that (1)-(4) below is true. (5) is also true—see the argument for (i).

⁶ Perfectionist goods are not equivalent to, and do not necessarily correlate with, prudential goods. See L. W. Sumner, *Welfare, Happiness, and Ethics* (Oxford: Oxford University, 1996), pp. 23-4.

And the only way (1)-(5) could all be true is if (6) is true. Therefore, (6), which is just (ii), is true.

So 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.

And, as I argued above, the following is also true:

- (5) The reason the agent has to act self-interestedly is a nonmoral reason. (This follows from (i).⁷)

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

- (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. (This is (ii).)

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

⁷ I actually don't need (i) for my argument for (ii) to go through. All I need is the claim that in many typical cases of agent-centered options, the reason the agent has to perform the self-interested alternative is a nonmoral reason.

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.

3. How (i) and (ii) motivate DRAC.

If (i) and (ii) are correct, then the permissibility of our actions isn't simply a function of what requiring reasons (a requiring reason is just a reason with some moral requiring strength) we have. This is potentially problematic for consequentialism, for consequentialism can naturally be thought of as the view that the better an outcome, the more requiring reason we have to produce it.⁸ Admittedly, traditional forms of maximizing act-consequentialism hold this view, but if (i) and (ii) are correct, then the act-consequentialist will have to make some adjustments. The adjustments needed are just the sort we find in SOU and other versions of DRAC. First, given (ii), given that morality is a function of both what we have best requiring reason to do and what we have best reason, all things considered, to do, the consequentialist will need one ranking corresponding to the former and one ranking corresponding to the latter. So if the consequentialist is one who thinks that the better an outcome the better reason there is to produce it, then the consequentialist will need a ranking of outcomes in terms of their moral value and a ranking of outcomes in terms of their all-things-considered value. Second, given (i), given that agents have no requiring reason to promote their own self-interest but only a requiring reason to promote the interests of others, we should interpret

⁸ See Howard-Snyder (1994).

SOU's first auxiliary ranking, the one that ranks outcomes in terms of total utility of everyone, as a ranking of outcomes in terms of their all-things-considered value and interpret SOU's second auxiliary ranking, the one that ranks outcomes in terms of total utility for everyone but the agent, as a ranking of outcomes in terms of their moral value. (Keep in mind that whereas moral value gives rise to requiring reason, non-moral value gives rise to only non-requiring reasons.)

So SOU's dual-ranking structure is not *ad hoc* at all. The move to a dual-ranking structure is not simply a move that is made to accommodate the self-other asymmetry. It is a move that is foisted upon the consequentialist given (i) and (ii). It might be objected, though, that my arguments for (i) and (ii) presupposes that there are agent-centered options of the sort that commonsense morality takes there to be, the sort of which (1)-(4) are true. But I don't see why the consequentialist can't help herself to (i) and (ii) given that her main rivals all presume that there are such agent-centered options and so are committed to the truth of (i) and (ii). So these non-consequentialist rivals can't object that (i) and (ii) aren't true. And if they are true, then the consequentialist has a principled reason for making the move to a dual-ranking structure.

V. What's that function?

A. The Function

If the permissibility of our actions isn't simply a function of what moral reasons we have but some function of both moral reasons and nonmoral reasons, what is that function. Call a reason that has some moral requiring strength a *requiring reason*. I'll defend the following meta-criterion:

MP S's performing *x* is *morally permissible* if and only if there is no available alternative that S has better requiring reason to perform and no worse reason, all things considered, to perform.

B. In Defense of MP

I'll defend this meta-criterion by arguing for the following three claims about the relationship between moral permissibility and reasons for action (both requiring and non-requiring), the conjunction of which entails MP:

CI It is always morally permissible to do what one has best requiring reason to do.

- CII** It is always morally permissible to do what there is decisive reason to do, all things considered. (There is decisive reason to do x if and only if there is better reason to perform x than to perform any other available act alternative.)
- CIII** If it is morally permissible to do x and there is better requiring reason to do y , then it is morally permissible to do y provided that there is no available alternative (to y) that one has even better requiring reason to do and no worse reason to do, all things considered.

C. In Defense of CI-CIII

CI: CI is hardly controversial, but, in any case, it follows from the definition of a requiring reason. By definition, only requiring reasons are capable of generating moral requirements. So even if an agent has better non-requiring (moral or non-moral) reasons to do other than that which she has best requiring reason to do, these reasons can't generate a moral requirement to do otherwise. And if the agent isn't morally required to do other than what she has best requiring reason to do, it follows that it is permissible for her to do what she has best requiring reason to do.

CII: CII is, by contrast, considerably more controversial, but only if we are unwilling to concede that non-requiring reasons can be justifying reasons, making it permissible to perform an act that would otherwise be impermissible. Once we concede this (and I've argued that we must—briefly here and in detail elsewhere), there is no more plausible claim about when non-requiring reasons make it permissible to perform an act that would otherwise be impermissible. For if the non-requiring reasons in favor of performing a given act don't make that act morally permissible when they defeat all other reasons (requiring and non-requiring), then when? Therefore, we should accept CII, i.e., that it is always morally permissible to do what one has decisive reason to do, all things considered. Besides, CII allows us to preserve the idea that morality has a certain rational authority such that doing what one is morally required to do is always what one has decisive reason to do, all things considered.

CII is strong moral rationalism. CII says: If S has decisive reason to do x , then S is morally permitted to do x . Now let ' $\sim x$ ' stand for 'something other than x '. Note that ' S has decisive reason to do x ' is equivalent to 'It is not the case that S has decisive reason to do $\sim x$ '. And note that ' S is morally permitted to do x ' is equivalent to 'It is not the case that S is morally required to do $\sim x$ '. Thus CII is equivalent to: If it is not the case that S has decisive reason to do $\sim x$, then it is not the case that S is morally required to do $\sim x$. By contraposition, this is

equivalent to: If S is morally required to do $\sim x$, then S has decisive reason to do $\sim x$. And, from this, we can generalize to: If S is morally required to ϕ , then S has decisive reason to ϕ . And this is just strong moral rationalism.

So the thought behind CII is that if non-moral reasons can prevent moral reasons from generating moral requirements, then there is no more plausible principle as to when they do so than strong moral rationalism.

One seemingly plausible alternative to CII is the view that it is always morally permissible to do what there is sufficient reason to do (all things considered), and so permissible to do what there is sufficient reason to do even when there is better requiring reason to do something else. But this isn't as plausible as CII, for this would allow one to do something that one has less requiring reason to do even when there is just as much reason to do what there is more requiring reason to do—I'll explain why this is implausible in the course of defending CIII below, where I'll be giving an example where one has more requiring reason to save two children as to save one child and no less reason, all things considered, to save the two children than to save one child. The only remaining alternative to CII is the view that it is sometimes morally permissible to do what there is less requiring reason to do even when there is insufficient reason, all things considered, to do so. But this is clearly implausible.

CIII: CIII expresses the sensible presumption that if it is morally permissible for P to perform x even though P has better requiring reason to perform y (x 's being permissible because that is what P has better, albeit non-requiring, reason to do), then it is also morally permissible for P to perform y . After all, if P has better requiring reason to perform y than to perform x , then it would be morally better if P were to perform y instead of x . The only instances in which this presumption of permissibility will ultimately prove false is where P has even better requiring reason and no worse reason, all things considered, to do something else, say, z . In that case, there would be no good reason (i.e., no undefeated reason) why P should not choose to perform z instead of y and one good reason (i.e., an undefeated requiring reason) why P should choose to perform z instead of y . Thus, if P is not going to perform x , P should, then, perform z , not y . And so we should accept CIII: it is morally permissible to choose to do what one has better requiring reason to do over what one has worse requiring reason to do if and only if there is no available alternative that one has even better requiring reason to do and no worse reason, all things considered, to do.

To illustrate, suppose that there is a burning building in which twelve children are trapped. Assume that I've called the fire department, but it is obvious that most, if not all, the children will perish before they can arrive unless I act now

and rescue the children myself. Lacking the proper fire-fighting equipment, there is no way for me to rescue the children without suffering severe injuries in the process, injuries that will require me to be hospitalized for many months and that will result in my permanent disfigurement. Let's suppose that, given these great costs, what I have best reason to do, all things considered, is to refrain from making any rescue attempts. But let's suppose that given the self-other asymmetry and given that what's at stake for me (my suffering and permanent disfigurement) is less than what's at stake for the children (their lives), what I have best requiring reason to do is to rescue as many children as possible. Assume, though, that the more children I rescue, the more times I'll need to enter the burning building and thus the more severe my resulting injuries will be. Now, let's suppose: (1) that a_5 is the act of waiting for the fire fighters to arrive, rescuing no children, (2) that a_6 is the act of entering the burning building once and rescuing two children, and (3) that a_7 is the act of entering the burning building once and rescuing a relatively unimportant file cabinet.

VI. Commonsense Consequentialism (CC)

When we take the consequentialist idea that the reasons for action are a function of what available worlds we have reason to prefer over others and combine that with the notion that moral status is a function of both moral reasons and nonmoral reason and that function is MP, we get the following version of act consequentialism:

Commonsense Consequentialism: S 's performing x is morally permissible if and only if there is no available act alternative that would actualize a possible world that S has both (i) more moral reason to want to be actualized than to want Wx to be actualized and (ii) at least as much reason, all things considered, to want to be actualized as to want Wx to be actualized. (Let Wx be the possible world that would be actualized were S to perform x .)

VII. Splawn's Objection

Consider the case below, where an agent named Jill has the following three mutually exclusive and jointly exhaustive options:

act	$U(x)$	$U_s(x)$	$\underline{U}_{-s}(x)$	$\underline{U}_{+s}(x)$	moral status
a_8	-40	10	-50	50	permissible
a_9	90	-10	100	0	permissible
a_{10}	44	-5	49	-1	impermissible

On SU, a_9 outranks a_{10} , because both (1) $U_{-s}(o_9) = 100 > U_{-s}(o_{10}) = 49$ and (2) $U_{+s}(o_9) = 0 \geq U_{+s}(o_{10}) = -1$. And so a_{10} is, on SU, impermissible. Yet a_{10} is considerably better for others than a_8 (a permissible alternative) is: $U_{-s}(o_{10}) = 49 > U_{-s}(o_8) = -50$. So why, Splawn might ask, is it permissible for S to perform a_8 but not a_{10} ? Sure, it's more costly for S to perform a_{10} than it is for S to perform a_8 , but Splawn would argue that we should accept the following plausible-sounding principle: "it is always permissible for an agent to make a self-sacrifice so long as that sacrifice does not bring about worse consequences for others" (2001, p. 330). Splawn would also argue that any theory that fails to accommodate this principle fails to accommodate the self-other asymmetry. He says, "[I]sn't that *exactly* the suggestion of the self-other asymmetry that sacrificing one's own good for the benefit of others is permissible? And isn't that *exactly* what [a_{10}] is: a considerable self-sacrifice on the part of S in order to bring about some good for others?" (2001, p. 330). I will argue, however, that this plausible-sounding principle is, in fact, *implausible* and that a theory doesn't need to accommodate this principle in order to accommodate the self-other asymmetry.

To see why the principle is implausible, we'll need to get clear on what precisely it says, specifically, what we're supposed to be comparing when we judge that a self-sacrificing act does or does not "bring about worse consequences for others." One suggestion might be that we're comparing what the total utility for others would be were the agent to interfere with ongoing causal processes by performing the self-sacrificing act with what it would be were the agent to refrain from interfering with ongoing causal processes and omit the self-sacrificing act. But I don't think that this is what Splawn has in mind, and, in any case, the principle is clearly implausible on such an interpretation. Consider again the burning building case from above and suppose that I take only one trip into the burning building and rescue the file cabinet instead of the two children. Did I act wrongly? Well, had I not performed this self-sacrificing act, the consequences would not have been worse for others (on this interpretation), for had I done nothing and just waited for the fire fighters to arrive, the two children would have still died. So, on this interpretation of the principle, taking one trip into the burning building and rescuing the file cabinet instead of the two children is permissible, which is absurd.

A more plausible suggestion, therefore, might be that we are, for each available alternative, to compare what the total utility for others would be were the agent to perform that alternative with what it would be were the agent to perform the self-sacrificing act in question. And so long as there is at least one alternative where others would be no worse off, then the self-sacrificing act is deemed permissible. On this interpretation of the principle, it is permissible for Jill to perform a_{10} (as Splawn thinks it should be), for the consequences for others would be no worse were Jill to perform a_{10} as opposed to a_8 : $U_{-s}(o_{10}) = 49 > U_{-s}(o_8) = -50$; indeed, the consequences would be better for others. The problem, though, with this interpretation is that it ignores the fact that a_9 is a permissible alternative. And although it does seem

permissible for Jill to perform a_{10} were a_8 and a_{10} the only available options, it is arguably impermissible for Jill to perform a_{10} given the availability of a_9 as an option. It's true that Jill does better by others by performing a_{10} instead of a_8 , but she can do even better by others by performing a_9 . And what justification does she have for not doing even better by others by performing a_9 ? Given that Jill has more reason, all things considered, to perform a_9 than to perform a_{10} , there would seem to be no good reason why Jill shouldn't perform a_9 instead of a_{10} . Thus Jill has no excuse for not doing more for the sake of others by performing a_9 instead of a_{10} . Sure, a_9 is more costly, but, by hypothesis, we're supposing the reason that Jill has to avoid this extra cost is insufficient to defeat the reason Jill has to do a_9 —namely, that a_9 will do more good for others. So although it is often permissible to make a self-sacrifice so as to benefit others, it is not permissible to do so when one could have benefited others even more without making any greater of a self-sacrifice than one has reason to make.

So when Splawn asks, "What possible justification could be given for saying [a_{10}] is wrong?" (330), the answer is that there is better requiring reason and no worse reason, all things considered, to perform a_9 instead. In choosing a_{10} over a_9 , Jill would be unreasonably selfish. Jill would be benefiting herself by 5 utiles at a cost of 51 utiles for others. So even if we grant, as SU does, that Jill has ten times more reason to benefit herself than to benefit others, we must still concede that Jill has more reason to benefit others by 51 utiles than she has to benefit herself by 5 utiles.

What's more, we find that, on this interpretation, Splawn's plausible-sounding principle falls victim to the same counter-example that it did on the first interpretation. On this second interpretation, it's also permissible to enter the burning building and save the file cabinet instead of the two children, for waiting for the fire department is an available option that is no worse for others than entering the burning building and rescuing the file cabinet is.

Indeed, the only way to interpret Splawn's principle so that it doesn't get this implausible result is to hold that we are, for each available alternative, to compare what the total utility for others would be were the agent to perform that alternative with what it would be were the agent to perform the self-sacrificing act, and then to hold that if there is even one alternative where others would be worse off, then the self-sacrificing act is *impermissible*. But, on this interpretation, a_{10} would, contrary to what Splawn says, be impermissible, for others are worse off on a_{10} than on a_9 . Since Splawn insists that the principle implies that a_{10} is permissible, this can't be how Splawn wants us to interpret this principle. So we must accept one of the first two interpretations of the principle. But, in that case, the principle is quite contrary to commonsense, implying that it is permissible to enter the burning building once and save the file cabinet instead of the two children. We should not, then, conclude that SU is objectionable given that it violates Splawn's principle. Nor should we conclude that SU fails to incorporate the self-other asymmetry in failing to accommodate Splawn's principle.