Consequentialism, Reasons, Benefits, and the Good

David McNaughton and Piers Rawling

We explore the relationships among consequentialism, practical reasons, benefits and the good. While, in the end, we reject consequentialism, en route to that conclusion we'll argue that, partly by paying careful attention to such matters as the distinction between benefits and the good, the view can cover more of the practical territory than some of its opponents suppose. Failure to keep in mind that distinction has bedeviled a number of classic discussions, including Ross's defense of deontology, Butler's discussion of divine utilitarianism, and Parfit's Repugnant Conclusion.

I

Reasons

Donald Davidson (1980) famously argues that practical reasons are mental states with causal powers. But we adopt a different usage, according to which practical reasons are facts – the fact that it's cold outside, say, is a reason for you to wear your coat.

It's important to note that there are two facts lurking here. Your reason is the first: it's cold. But there's also a second: the fact that the first fact is a reason. We have, then, a two-tier view of practical reasons. At tier one are the reasons; at tier two are the facts that the tier one facts *are* reasons. Experience tells us that it's easy to muddle this distinction, so perhaps it helps to appreciate that the two tiers give rise to different possibilities of error: you might be mistaken about the weather (tier one error); or you might fail to realize that cold weather is reason to wear a coat (tier two error).

In some contexts 'reason' refers only to a cause, as in, 'metal fatigue was the main reason for the bridge's collapse'. For Davidson, 'reason' refers, in psychological contexts, to a cause that rationalizes – if, say, his reason for calling was that he wanted to persuade her to go, and he believed he could do so by calling her, then, according to Davidson, these mental states caused his calling. But they also rationalize his behaviour. And rationality, unlike causality, is a normative notion: ceteris paribus, we *should* be rational. Our notion of a reason is also normative: if it's cold, and this fact is a reason for you to wear your coat, then (roughly speaking), in the absence of countervailing reasons, you should wear your coat.

Rationality and reason (in its normative senses) are part of a broader family of intertwined normative notions. What is normativity? Like Parfit, we do not think it can be analyzed in other terms; all we can do is give illustrations of normative notions, and their relations. On our usage, for example, a normative requirement is not made so by the fact that most people *do* heed it (if they do), but by the fact that they *should* heed it – they have overriding *reason* to heed it. Or consider another pair of normative notions: the evaluative and the moral. To say that a potential act of yours would make the world better is, on Scanlon's (1998) 'buck-passing' account of the good, to say that you have a reason of a certain kind to perform it. If an act would be morally wrong, then (on many views, at least – we'll discuss a potential counterinstance below) you have overriding reason not to perform it. The list of other normative notions includes harm, benefit, innocence, desert, justice, fidelity, gratitude, reparation; and these too are inter-related. W.D. Ross (1930, p. 21), for example, takes justice to be concerned with the distribution of benefits and harms in accord with desert; and desert, of course, is tied to innocence. The normative is also tied to the non-normative. Supervenience holds, for instance: innocence, say, cannot be lost without a change in non-normative circumstances. And there may be non-normative items on lists of benefits, harms, and valuable (or not) features of states of affairs. Not all such items are non-normative, however; and normativity also enters when we ask what unifies the lists.

According to what we'll call the "simple view" of reasons, normative notions do not enter into their content. On this view practical reasons (tier one facts) are non-normative facts. It is tier two facts that are normative – i.e., have normative content.[1] A practical reason (on this view) is not itself a normative fact; it is the fact that you have it that's normative.

But is this view correct? Suppose an act you're contemplating would cause someone else undeserved harm. This, we think, is a reason against performing it, the normativity of the notions of harm and desert notwithstanding. We reject, then, the simple view of reasons, and allow that there are normative tier one facts – i.e., what we call "normative reasons" (usage varies).[2]

---

[1] Note that adverting to content is crucial here – normative facts can be *picked out* using non-normative vocabulary, at least on occasion: for instance, 'the first fact David thought of when waking today' might pick out the fact that he has reason to go into the office.

[2] Rejecting the simple view means abandoning a simple account of the supervenience of practical reasons on the non-normative: on the simple view, there cannot be a change in what you have reason to do without a change in the non-normative facts because your reasons are among the non-normative facts. If we allow normative facts at tier one, this, obviously, no longer holds. However, supervenience still holds provided that, to take the current case as an illustration,

So much for our brief remarks on the content of reasons; what about their strength? Suppose you have a sum of money that you can either donate to Oxfam or use to take your family out to dinner. What do you have more reason to do? On the one hand, the donation to Oxfam would do some good; on the other, going out to dinner would (to use a philosopher's catch-all) benefit you and your family. The donation to Oxfam would do good (at least in part) by benefitting people, of course; and your taking the family to dinner might do some good in addition to its benefits. But are there other primitives at the level of generality of benefits and goodness to be considered as contributors to reason-strength? Their opposites – harms and badness – play in, of course. And of harms and benefits it can asked: harmful or beneficial to whom or what? But that, to us, appears to be it – so here's a piece of advice: if it doesn't benefit anyone (or any thing, perhaps) or do any good, don't do it, because you have no reason to.

The next issue is the relation between benefits and value – are benefits always good? Given that they are, at least, sometimes so, do they contribute to reason strength only via their contribution to the good? To these and related topics we now turn.

II

Benefits, Value, and Reason Strength

In acting you modify the state of the world. Some states are better than others; and some states are better *for you* than others. The 'better than' relation ranks states in accord with their goodness or value; the 'better for x than' relation ranks states in accord with how beneficial they are to x. Terminology can be confusing here: some authors distinguish between impersonal and personal value, but as we use the term, value is always impersonal; it is benefits that are personal. Each is measured on an objective scale – it is not the case, for example, that x's benefit scale varies in accord with the perspective of the evaluator. And care must be taken not to conflate benefit scales with the value scale.

Consider, for instance, the following case from Ross (1930, pp.34-5):

---

harm and desert have non-normative supervenience bases – that is, there cannot be two cases that differ in terms of harm or desert without some non-normative difference. This seems to us obviously true.

> Suppose … that the fulfillment of a promise to A would produce 1,000 units of good for him, but that by doing some other act I could produce 1,001 units of good for B, to whom I have made no promise, the other consequences of the two acts being of equal value …

Ross sees this case as a potential counterexample to the view that right acts are 'those productive of the best possible consequences' (ibid. p.34) – i.e., consequentialism. The consequentialist, however, might reply by distinguishing between benefits and the good. Ross seems to suppose that providing x units of benefit to someone (i.e., x units of good *for* them) produces x units of good simpliciter. But the consequentialist can deny this and claim that units of benefit are independent of units of good. She might, for instance, hold the following view of the example: keeping the promise to A has value x; providing 1,000 units of benefit to A has value y; providing 1001 units of benefit to B has value z. She then asks which is greater: x+y or z? And the answer is determined by the particulars of the case.[3]

G. E. Moore is one of Ross's targets. Moore's theory, Ross claims, 'says, in effect, that the only morally significant relation in which my neighbours stand to me is that of being the possible beneficiaries of my action' (ibid. p.19). We're not sure that's the case. But even if it were, it does not follow that Moore has to equate the value of a state with the sum of the benefits it includes.

We are not consequentialists, but we think that it is harder to dismiss than some of its opponents suppose. According to what we'll call "simple consequentialism", you only have reason to phi if doing so will produce a better state of the world than doing nothing (where the value produced by an act includes the value of the act itself); and the (overall) strength of your reason(s) to phi varies in accord with the value of phi-ing, so that what you have most reason to do is perform the act that will maximize value (or one of them in the case of a tie). The content of your reasons, however, need not mention value – the simple consequentialist could adopt something akin to Scanlon's (1998) 'buck-passing' view (mentioned above), and claim that, although value determines reason strength, it need not enter reason content. Your reason to put on your coat is that it's cold; the strength of this reason is a matter of the value of the act relative to other alternatives available to you. But what determines the value of an act?

---

[3] Ross's oversight is made explicit in his discussion (ibid. p.35) of a variant of the above example, when he starts by speaking of a disparity in the provision of 'units of good for A', and ends by speaking of 'a disparity of good' simpliciter.

The (normative) hedonist holds, perhaps, the simplest view: pleasure is the only good (and pain the only bad). This might be construed as denying that benefits are relevant to value. Alternatively, however, the claim that pleasure is the only good can be seen as emerging from the conjunction of: (1) pleasure is the only thing beneficial; and (2) benefits are the only good. In addition, on this approach, the hedonist holds that (3) pleasure is always beneficial, and benefits are always good.

The simple consequentialist need not, of course, be a hedonist. The latter has no room for the thought that one distribution of benefits is better than another; the simple consequentialist, on the other hand, can incorporate, say, the view that justice – in the sense of distributing benefits in accord with desert – is itself good (see, for instance, Ross, ibid. pp.26-27, 138). This adds to the list of goods, of course, in denial of (2); and (3) fails also, since a benefit going to someone who deserves harm, for example, is bad. Note that in order even to *raise* the issue of whether one distribution is better than another, there must be something to be distributed – namely, benefits (and harms).

Since Ross endorses a prima facie duty of justice, which regards the distribution of benefits in accord with desert as itself a good, it is all the more surprising that, he conflates benefit and value in his argument (above) to show that consequentialism cannot accommodate the duty of fidelity. But this conflation (or confusion) even extends to his discussion of justice itself. Early in the chapter he allows that the prima facie duty of justice "comes under the general principle that we should produce as much good as possible" (p. 27); i.e. he allows that consequentialism can accommodate justice. Yet a few pages later, he tries to show, by appeal to the same style of argument that he used in the case of fidelity, that 'ideal utilitarianism' cannot accommodate justice.

> Suppose that A is a very good and B a bad man, should I … think it self-evidently right to produce 1,001 units of good for B rather than 1,000 for A? Surely not. I should be sensible of a *prima facie* duty of justice, i.e. of producing a distribution of goods in proportion to merit, which is not outweighed by such a slight disparity in the total goods to be produced (p. 35)

Ross is by no means alone in failing to keep in mind the distinction between value and benefit when discussing consequentialism. Since the error is both pervasive and insidious, a couple of further examples may reinforce its importance.

Joseph Butler, while insisting that *our* moral thinking is not utilitarian in nature, is nevertheless inclined to suppose that God is a utilitarian. Robert Louden ('Butler's Divine Utilitarianism', *History of Philosophy Quarterly*, 1995, pp. 265-80) argues that Butler offers a theological version of what Bernard Williams disparagingly called 'Government House Utilitarianism'. Given our ignorance of the wider implications of our actions, a wise Governor will instill in us a respect for veracity and justice in themselves, and a propensity to favour friends and family over strangers, because such traits are most conducive to the general welfare. Louden asks why Butler supposes that God's moral character can be exhaustively characterized as benevolent, understood as a disposition to promote the happiness of others, rather than also attributing to Him "deontological traits such as justice and truth-telling" (275). This is a false dichotomy, since it omits pluralist consequentialism. Butler himself leaves no room for this position, because he repeatedly identifies a concern for the *good* of the world, or of creation, with a concern for the *happiness* of its creatures. It is striking that Louden follows Butler in this assumption. Since Butler insists that God is a righteous Governor he is driven to the implausible claim that God may punish the wicked retributively and yet be "simply and absolutely benevolent" (AR I iii 3). But such punishment can, of course, be consistent with a concern for the *good* of the world.

Parfit's discussion of his 'Repugnant Conclusion' and related topics is long and subtle, and there is much that we shall not address. But here are the basics that are relevant for our present purpose. Consider first Parfit's 'Impersonal Total Principle':

> If other things are equal, the best outcome is the one in which there would be the greatest quantity of whatever makes life worth living.[4]

This principle supposes that benefits are the only thing of value. And here's a well-known consequence of it:

> *The Repugnant Conclusion.* For any possible population of at least ten billion people, all with a very high quality of life, there must be some much larger imaginable population whose existence, if other things are equal, would be better, even though its members have lives that are barely worth living.[5]

However, there does seem to be something wrong with a view of value that entails the conclusion that the repugnant world is best. Our view of value and

---

[4] Ibid., 387.

[5] Ibid., 388.

benefits avoids it. How? By denying that benefits are the only things of value. We might be accused of simply denying this in order to avoid the repugnant conclusion, with no further rationale. But we have already seen good reason to reject this view. And a further argument for our claim is provided by the repugnant conclusion itself. Why is it repugnant? Well, clearly the repugnant world is worse than the first alternative in Parfit's conclusion (we can do no more than appeal to widely shared intuitions here), even though the sum total of the benefits within the repugnant world is the larger. Which all goes to show that value is not to be measured by the sum total of benefits.

Setting aside (1), the issue of whether pleasure is the only thing beneficial, let's return to the issue of reason strength. According to simple consequentialism, you only have reason to pursue a benefit (for yourself or others) if doing so will increase the good; and the strength of your reason so to act is proportional to the amount of good you would produce. You might, say, have most reason to pursue your own benefit in some circumstance, on this view, but only if you would thereby produce the best (in the sense of maximizing value) state. Benefits are relevant to reasons here, but only indirectly: while benefits are relevant to the value of a state, the strength of your reason(s) to produce that state is proportional only to its value.

There are, of course, other possibilities. According to the (normative) egoist, for instance, value is irrelevant to reason strength – rather, the strength of your reason to do something is a matter only of how much it would benefit you. We try to occupy an intermediate position between these two extremes. Like the advocate of egoism, we see benefits as playing a direct role in our practical reasons. And like simple consequentialism, we also see the good as playing such a role.

Here's one component of our position: you have reason to perform some act only if, in comparison to doing nothing, either so acting will give rise to benefit (for someone or something) or good, or both. If your contemplated act is of no benefit or good (relative to doing nothing), you have no reason to perform it.[6] And, *pace*

[6] Two aspects of this claim might initially appear puzzling. First, why the comparison to doing nothing? This is to accommodate cases in which, if you act either harm or badness will result, but if you do nothing even more harm will arise, or things will go even worse. We do not want to rule out your having reason to act in such unfortunate circumstances. (Implicit in our view, then, is the thought that reducing harm counts as producing benefit; and reducing badness counts as producing good.) Of course, the notion of 'doing nothing' is tricky, and we certainly don't want to enter the debates about acts and omissions. What it would be to do nothing will, however (we hope), be clear in any given case. Second, why not strengthen the claim from 'only if' to 'if and only if'? Well, suppose you could benefit a justly imprisoned felon by helping him escape. His

simple consequentialism, we think that there are occasions on which you have more reason to perform an act that is worse (in terms of value) than some other; and you do not always have most reason to maximize the good.

On our view, the strength of your reasons to benefit yourself or your special others (those to whom you have special ties, such as your friends and family, and those to whom you have made binding promises etc.) may exceed whatever contribution the conferral of such benefits would make to the general good. You may have reason to benefit a stranger, but, other things equal, the strength of such a reason, we claim, is a function of how much good would be accomplished. The thought is, roughly speaking, that when benefitting strangers, you should distribute the benefits so as to do the most good. When it comes to yourself and your special others, however, we contend that there can be more reason to perform beneficial acts than their (the acts') value warrants. We are proposing, then, that although the strength of our reasons to benefit ourselves and our special others does not vary with value alone (for brevity, we sum this up by saying that we have "personal" and "special" reasons, respectively), yet when it comes to benefiting or harming those outside this circle, value is the sole measure of reason strength.

Our views on personal and special reasons, then, set us apart from the simple consequentialist. We see her as confronting at least a *prima facie* dilemma. Either she acknowledges the existence of benefits or she doesn't. On the latter horn, she is in the position of having to deny even the possibility of a debate over distributional concerns (see above). On the former horn, she has to counteract the plausible thought that, on occasion at least, we have personal and special reasons the strength of which is disproportionate to the value of the acts in question. For instance, a state in which I benefit may well be equally valuable to one in which you do. Let us suppose it so. Then, as far as value goes, you have as much reason to benefit me as you do to benefit yourself. But does such parity of value rule out the possibility that you have more reason to benefit yourself than to benefit me? Or consider a case in which your receiving some benefit would make the world worse, might you not have some reason to pursue it? (Doesn't even a justly imprisoned felon have reason to escape?)

Are there other considerations that favour our view of personal and special reasons? We think there are. We'll focus here on the case of friendship – in

---

benefit notwithstanding, you may have no reason to. However, if an act would do some good (in comparison to doing nothing), then you do have some reason to perform it.

essence we agree with a version of what Jackson (1991) refers to as the ‘nearest and dearest objection’ to consequentialism.

III

The Nearest and Dearest Objection

Here is Jackson’s statement of the nearest and dearest objection:

> Our lives are given shape, meaning and value by what we hold dear, by those persons and life projects to which we are especially committed. This implies that when we act we must give a special place to those persons (typically our family and friends) and those projects. But, according to consequentialism classically conceived, the rightness and wrongness of an action is determined by the action’s consequences considered impartially, without reference to the agent whose actions they are consequences of. It is the nature of any particular consequence that matters, not the identity of the agent responsible for the consequence. It seems then that consequentialism is in conflict with what makes life worth living. I take this to be one part of Bernard Williams’s well-known attack on consequentialism. (Jackson 1991, p.461)

Jackson’s response to this objection (ibid. p.472ff) is to point out that it’s very difficult to know which of our potential acts will do good, and that ‘our opinions on the matter are much better founded in the case of people we know well precisely because we know them well’ (ibid. p.475). Thus ‘the good consequentialist should focus her attentions on securing the welfare of a relatively small number of people, herself included, not because she rates their welfare more highly than the welfare of others but because she is in a better position to secure their welfare’ (ibid. p.481).

Jackson is trying to avoid any appeal to a ‘biased value function’ (p.475) – that is, to the thought that our nearest and dearest are more valuable than those outside our intimate circle. But there is an alternative consequentialist response to the nearest and dearest objection that neither appeals to our epistemic limitations nor to a biased value function. The consequentialist can think that ‘deep personal attachments’ (Hooker, 2000, p.141) are themselves valuable – she might hold, say, that friends attending to friends is a good thing: the world is better for these caring relationships. If consequentialist Alan takes this line, he need have no

scruples about benefitting Betty because she is *his*[7] friend – not because he holds that Betty is herself better than others that he might benefit, but because he is promoting a friendship, and friendship is a good.

But what would Betty think of Alan's reasons to benefit her? As we've pointed out, if the consequentialist takes a buck-passing approach, value need not enter the content of Alan's reasons. If so Betty won't be able to complain that, say, Alan's reason for presenting her with flowers is not his love for her, but that love relationships are valuable. Nevertheless, if she looks to how Alan weighs his reasons, she might be less than pleased when she realizes that he will abandon her if the going is better elsewhere. A similar sort of problem confronts Jackson's epistemic solution – imagine Alan explaining to Betty that he bought flowers for her as opposed to Clarissa because he knows which variety she (Betty) likes, but is not sure of Clarissa's tastes.

These difficulties can be avoided, perhaps, if consequentialism goes indirect – the idea being, roughly, that agents are to be trained to do what is most valuable without consciously entertaining thoughts about value.[8] Alan's account of his reasons and how he weighs them would make no reference to value or epistemic limitations – that's outside his purview. We're dubious of this approach, however – for one thing, we see the ideal agent as being able, should she have the time and inclination, to uncover the nature and relative strengths of her practical reasons. One of the agreeable features of simple consequentialism, from our perspective, is its straightforwardness about practical reasons: it is no part of the theory that the truth about practical reasons be hidden from those it governs.

We have been arguing that fully informed consequentialists would have difficulty maintaining relationships because, in short, acknowledging that value is the sole determinant of the strength of their reasons would undermine the parties' claims that they regard one another as "special". Beyond this problem, however, there is the question of whether psychological attachment is all there is to friendship.

Consider loyalty between friends. It is not just that friends are psychologically attached to one another – spend time with each other, support each other, and so on. In addition, a friend expects your loyalty because she is your friend. If you

[7] Jackson's remark above about consequentialism's lack of reference to agents may be misleading.

[8] As Julia Driver noted in conversation, if the buck-passing approach to reason content is correct, then such agents can still consciously access the content of their reasons; it's reason strength that is to be, as it were, kept from them.

betray her, she has a complaint against you that no one else has. And this tie of loyalty between friends appears essential to friendship. Friends come through for one another; someone who neither came through for you, nor believed she should, would not be loyal and so would not be a friend.

So is loyalty a possibility for a thoroughgoing simple consequentialist? We think not. If it were, this would be because loyalty is good. But then the simple consequentialist would have, for example, more reason than not to be disloyal to her own friends if she could (in odd circumstances, admittedly) thereby bring about more loyalty in the world overall. And this would mean that she's not loyal to her friends after all.

Thus not only would a simple consequentialist have trouble maintaining personal attachments, but friendship would not even be a possibility on her position – so much the worse, perhaps, for attachments and friendship[9].

How do things look on our own view? As described above, we think that each of us has special reasons to benefit our friends, and this enables our view to accommodate loyalty and friendship. It is open to us to think that friendship is good; but, crucially, the strengths of your reasons to be loyal to, and otherwise benefit, your friends can exceed the value of doing so.

We are not, then, consequentialists. But neither are we deontologists – at least not if deontologists are taken to favour what have come to be known as 'constraints' (see, e.g., Alexander & Moore, 2008.), the topic of our next section.

IV

Constraints

A constraint is a prohibition against harming people, even in pursuit of good ends – even, indeed, to prevent a greater amount of the very kind of harm that is prohibited by the constraint in question.

For our purposes we first need an account of constraints in terms of reasons. While, in our view, there is no sharp division of practical reasons into the moral

[9] One of the authors heard Donald Regan take something like this line in a conference presentation.

and the non-moral, examples can be provided that are clearly on one side or the other. Your reason to choose a peach over an apple – that the former is sweeter – is non-moral. Your reason to give to Oxfam, on the other hand – that doing so will reduce innocent suffering – is moral. Or suppose that you promised to repay a debt on Thursday; this fact is a moral reason to do so. What about your reasons to favour your friends, discussed above? Some object to the idea that any of these are moral on the grounds that there is something less than ideal about doing things for friends out of a sense of obligation. But that is to confuse the issue of reason with that of motivation: it's quite possible to act on a moral reason – to refuse to betray a friend, say – out of affection.[10]

The moral reasons just mentioned fall within two categories: some are associated with promoting the good, and some with special ties. And some personal reasons may also be moral. But the advocate of constraints can be seen as claiming that there is a further category of moral reasons. Suppose, for instance, there were a constraint against killing the innocent. Then, whether this constraint be absolute or threshold, there would be at least one possible occasion on which the strength of your moral reason not to kill an innocent stranger would be greater than that which would correspond to the disvalue of the killing. On such occasions, your killing the innocent stranger would be bad, but not doing so would be worse (more innocents would be killed by others, say); yet the constraint would dictate that you have more moral reason not to kill. So the strength of your moral reason not to kill an innocent stranger does not vary only with the badness of doing so – call such a reason a “constraining reason”.

Our view is that the strength of your reason to perform a given act is a matter only of how much benefit or harm to yourself or those to whom you stand in special relations, and/or how much good or bad, would result. Thus we leave no room for constraining reasons. By assumption, constraints do not involve special relationships. And the badness of harm cannot do the job: the constraint *violator* in the previous paragraph would ensure that less bad comes about.[11]

---

[10] Setting aside ties, we tentatively suggest that you are morally required to phi iff:
(1) phi-ing is what you have most reason to do overall;
(2) the preponderance of your reasons to phi are moral; and
(3) you have more moral reason to phi than you have moral reason to do anything else.
And phi-ing is morally wrong iff there are moral reasons against phi-ing that outweigh the reasons in its favour.
Note that, unlike, say, Scanlon (1998), we do not think that an act’s being morally wrong is a *further* reason against it.

[11] Admittedly, violating a constraint might harm the violator (consider the psychological trauma, for example). But such harms to the violator do not provide her with *moral* reasons not to violate, and thus are not, presumably, what the advocate of constraints has in mind as grounding them.

Constraints, then, from our perspective, would require a further primitive concerning reason strength – they cannot be accommodated by appeal to benefit, harm or value. We see little prospect of a plausible rationale for including such an added extra. But arguing this case in detail here would take us too far afield.

V

Conclusion

Whereas, say, Bernard Williams (1979) explores the idea that among the necessary conditions of reasonhood are certain features of an agent's motivational psychology, we have looked in a different direction – at the claim that, roughly,

> (R) You have reason to perform some act only if benefit or good would ensue.

This is a very capacious claim, and all we have had the space to do here is defend an outline of our view as one of the contenders it accommodates. (We have said, for example, little about what is valuable or beneficial, nor about the details of weighing reasons.)

Even more capacious than (R), however, is our general framework: we see the practical landscape solely in terms of reasons. But, in conclusion, we note that this view is certainly not universally shared. Here's a quote from Singer (1999, pp. 308-309)[12]:

> [Nagel and I] were discussing 'Famine, affluence and morality', and Nagel was unable to accept that morality could be so demanding. But eventually it emerged that he was assuming that if morality did demand that we give so much to famine relief, then there must be overriding reason to do so. I was making no such assumption. On my view, I could recognize that if I were totally committed to doing what I ought to do, I would give away my wealth up to the point indicated in my article; but at the same time I may, without any irrationality, choose to be less than totally committed to doing what I ought to do. My own interests, or those of my family, may counteract the demands of morality to some degree, and I may think it reasonable to give in to them, while recognizing that it is morally wrong for me to do so. Once Nagel and I realized that we held these distinct

---

[12] Thanks to Doug Portmore for alerting us to this passage.

> understandings of morality, the practical difference between Nagel and myself over the demandingness of morality became less acute.

On Singer's view, then, morality and practical reasons are, to some extent at least, in different camps. Within our framework, by contrast, your moral reasons (if such there be) are thrown into the mix to be weighed along with all your other practical reasons. Singer sees morality, perhaps, as akin in some respects to the law – it is true, for example, that you do not always have overriding reason to obey the law of the land. Indeed, it may be asked whether you have any reason to do so, and Singer would seem to invite the general question, "Do I have reason to be moral?" Within our framework this question becomes, "Do I have reason to do what I have moral reason to do?" So for us the real question here is: "Are there any moral reasons?" – a different question from the first. For example, some Hobbesians might claim that, at least on occasion, we each have reason to be "moral" because it is in our self-interest. But they might deny that there are any moral reasons on the grounds that egoistic reasons to be moral are not moral reasons.

We, of course, think our framework is the one that reflects the truth about the relation between practical reasons and morality – otherwise morality fails to be action guiding in any direct sense. But we lack (thankfully) the space to elaborate on this further here.

References

Alexander, L. and Moore, M. (2008) ‘Deontological Ethics’, *The Stanford Encyclopedia of Philosophy (Fall 2008 Edition)*, E. N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2008/entries/ethics-deontological/>.

Butler, J. (1729) *Fifteen Sermons* London

Davidson, D. (1980) *Essays on Actions and Events* Oxford: Clarendon Press.

Harrison, R. (ed.) (1979) *Rational Action* Cambridge: Cambridge University Press

Hooker, B. (2000) *Ideal Code, Real World* Oxford: Clarendon Press.

Jackson, F. (1991) ‘Decision-Theoretic Consequentialism and the Nearest and Dearest Objection’ *Ethics* 101(3) 461-482

Jamieson, D. (ed.) (1999) *Singer and His Critics* Oxford: Blackwell

Moore, G.E. (1903/1966) *Principia Ethica* Cambridge: Cambridge University Press (page nos. refer to the 1966 paperback edition)

Ross, W.D. (1930) *The Right and the Good* Oxford: Clarendon Press.

Scanlon, T. (1998) *What We Owe to Each Other* Cambridge, Mass.: Harvard University Press.

Singer, P. (1999) ‘A Response’ in Jamieson (1999)

Williams, B.A.O. (1979) ‘Internal and External Reasons’ in Harrison (1979)