

Penultimate draft. For published version, see James Dreier (ed.) *Blackwell Contemporary Debates in Moral Theory*, 2006

Reasons without Demands: Rethinking Rightness

1. Introduction

My concern in this paper is to argue that consequentialist theories such as utilitarianism are best understood purely as theories of the comparative value of alternative actions, not as theories of right and wrong that demand, forbid, or permit the performance of certain actions. Consequentialist morality, I will argue, provides reasons for actions, without issuing demands (or permissions). Such an approach can answer the three related criticisms of consequentialism that it requires too much sacrifice of agents, leaves inadequate room for moral freedom, and does not allow for supererogation.¹ These criticisms focus on the maximizing feature of the most common forms of consequentialism, pointing out that maximization leaves little room for options. I will also argue that these criticisms have very little force against more traditional versions of consequentialism, on any reasonable understanding of what rightness amounts to. The rejection of rightness, though, does not address a different type of criticism of consequentialism. According to some, consequentialist theories are unacceptable, because they fail to account for constraints on permissible behavior. I will briefly discuss Peter Vallentyne's version of this criticism, which claims that such constraints are required to recognize what he (and others) calls the 'normative separateness of persons'. I will argue that, despite the undoubted rhetorical appeal of this phrase, it does not provide the basis for a convincing criticism of consequentialism. Either 'the normative

separateness of persons' signifies a feature fully accounted for by consequentialism (and probably not by rival theories), or it refers simply to the claims that (a) persons have certain rights, and (b) usually (or even always) rights not to be harmed are more stringent than rights to be aided. I shall conduct most of my discussion in terms of utilitarianism, since this is the most popular form of consequentialism. None of my points, however, will rely on the utilitarian value theory². I will also not devote much time to explaining the basic structure of utilitarianism, since both William Shaw and Peter Vallentyne do an excellent job in that regard.

2. The Demands of Utilitarianism.

The three criticisms of utilitarianism, that it requires too much sacrifice of agents, leaves inadequate room for moral freedom, and does not allow for supererogation, can be seen as applications of the more general criticism that utilitarianism is *too demanding*. But how, exactly, are we to take this criticism? Utilitarianism is too demanding *for what*? If I take up a hobby, say mountain climbing, I may well decide that it is too demanding *for me*. By that, I mean that I am simply not willing to accept the demands of this hobby. I may, therefore, decide to adopt the less demanding hobby of reading about mountain climbing instead. However, unless we adopt a radically subjectivist view of the nature of morality, according to which I am free simply to pick whichever moral theory pleases me, this approach will not work for the claim that utilitarianism is too demanding. When critics object to what they see as utilitarianism's demands, they are not simply declaring themselves unwilling to accept these demands, but are claiming that morality doesn't, in fact, make such demands. We are not, they claim, actually required to sacrifice our own

Reasons without Demands: Rethinking Rightness

interests for the good of others, at least not as much as utilitarianism tells us.

Furthermore, we really do have a (fairly) wide range of moral freedom, and there really are times when we can go above and beyond the call of duty. Since utilitarianism seems to deny these claims, it must be rejected.

How should a utilitarian respond to this line of criticism? One perfectly respectable response is simply to deny the claims at the heart of it. We might insist that morality really is very demanding, in precisely the way utilitarianism says it is. But doesn't this fly in the face of common sense? Well, perhaps it does, but so what? Until relatively recently, moral 'common-sense' viewed women as having an inferior moral status to men, and some races as having an inferior status to others. These judgments were not restricted to the philosophically unsophisticated. Such illustrious philosophers as Aristotle and Hume accepted positions of this nature. Many utilitarians (myself included) believe that the interests of sentient nonhuman animals should be given equal consideration in moral decisions with the interests of humans. This claim certainly conflicts with the 'common-sense' of many (probably most) humans, and many (perhaps most) philosophers. It should not, on that account alone, be rejected. Indeed, very few philosophers base their rejection of a principle of equal consideration for nonhuman animals merely on its conflict with 'common-sense'. Furthermore, it is worth noting that the main contemporary alternative to a (roughly) consequentialist approach to morality is often referred to as "common-sense morality"³. Those who employ this phrase do not intend the label itself to constitute an argument against consequentialism.

As I said, a perfectly respectable utilitarian response to the criticism that utilitarianism is too demanding is simply to insist that morality really is very demanding.

Reasons without Demands: Rethinking Rightness

However, there are powerful reasons to take a different approach altogether. Instead of either maintaining the demands of maximizing utilitarianism, or altering the theory to modify its demands, we should reject the notion that morality issues demands at all. In order to see why this might be an attractive option, I will briefly examine the alleged category of supererogatory actions, and an attempted modification of utilitarianism to accommodate it.

Maximizing utilitarianism, since it classifies as wrong all acts that fail to maximize, leaves no room for supererogation. A supererogatory act is generally characterized as an act which is not required, but which is in some way better than the alternatives. E.g. a doctor, who hears of an epidemic in another town may choose to go to the assistance of the people who are suffering there, although in doing so he will be putting himself at great risk.⁴ Such an action is not morally required of the doctor, but it produces more utility than the morally permissible alternative of remaining in his home town. The category of the supererogatory embodies two connected intuitions that are at odds with maximizing utilitarianism. First, it seems that people sometimes go beyond the call of duty. Maximizing utilitarianism would not allow that. To do your duty is to do the best thing you can possibly do. And second, people who fail to make certain extreme sacrifices for the greater good are usually not wrong. It seems harsh to demand or expect that the doctor sacrifice his life for the villagers.

The utilitarian can avoid these consequences by retreating to a form of satisficing utilitarianism.⁵ For example, one can allow that the boundary between right and wrong can in some cases be located on the scale at some point short of the best. This would allow that an agent can do her duty without performing the best action available to her,

Reasons without Demands: Rethinking Rightness

and it would make it possible for her to go beyond the call of duty. The position of the boundary between right and wrong may be affected by such factors as how much self-sacrifice is required of the agent by the various options, and how much utility or disutility they will produce. For example, it may be perfectly permissible for the doctor to stay at home, even though the best option would have been to go and help with the epidemic. On the other hand, if all the doctor could do and needed to do to save the villagers were to send a box of tablets or a textbook on diseases, then he would be required to do all he could to save them.

Satisficing versions of utilitarianism, no less than the traditional ones, assume that the rightness of an action is an all-or-nothing property. If an action does not produce at least the required amount of good, then it is wrong; otherwise it is right. On a maximizing theory the required amount is the most good available. On a non-maximizing theory what is required may be less than the best. Both forms of utilitarianism share the view that a moral miss is as good as a mile. If you don't produce as much good as is required, then you do something wrong, and that's all there is to it.

Utilitarianism has traditionally been viewed as a theory of right action. Utilitarians have employed theories of value, theories that tell us what things are good and bad, in functions that tell us what actions are right and wrong. The most common function from the good to the right is the maximizing one: an act is right if and only if it produces at least as much good as any alternative available to the agent, otherwise it is wrong. According to this maximizing function, rightness and wrongness are not matters of degree. Utilitarians are not alone on this score. Deontologists concur that rightness and wrongness are not matters of degree. There is an important difference, though. In

typical deontological theories, properties that make an action right and wrong -- e.g., being a keeping of a binding promise, a killing of an innocent person, or a telling of a lie -- are *not* naturally thought of as matters of degree. So one wouldn't expect the rightness or wrongness of an act to be a matter of degree for deontology⁶. But this is not the case with utilitarianism. Goodness and badness are clearly matters of degree. So the property of an act that makes it right or wrong -- how much good it produces relative to available alternatives -- *is* naturally thought of as a matter of degree. Why, then, is rightness and wrongness not a matter of degree?

3. Scalar⁷ Utilitarianism.

Here's an argument for the view that rightness and wrongness isn't an all-or-nothing affair. Suppose that we have some obligations of beneficence, e.g. the wealthy are required to give up a minimal proportion of their incomes for the support of the poor and hungry. (Most people, including deontologists such as Kant and Ross, would accept this.) Suppose Jones is obligated to give 10% of his income to charity. The difference between giving 8% and 9% is the same, in some obvious physical sense, as the difference between giving 9% and 10%, or between giving 11% and 12%. Such similarities should be reflected in moral similarities. A moral theory which says that there is a *really significant* moral difference between giving 9% and 10%, but *not* between giving 11% and 12%, looks misguided. At least, no utilitarian should accept this. She will be equally concerned about the difference between giving 11% and 12% as the difference between giving 9% and 10%. To see this, suppose that Jones were torn between giving 11% and

Reasons without Demands: Rethinking Rightness

12% and that Smith were torn between giving 9% and 10%. The utilitarian will tell you to spend the same amount of time persuading each to give the larger sum, assuming that other things are equal. This is because she is concerned with certain sorts of consequences, in this case, with getting money to people who need it. An extra \$5,000 from Jones (who has already given 11%) would satisfy this goal as well as an extra \$5,000 from Smith (who has given 9%). It does not matter whether the \$5,000 comes from one who has already given 11% or from one who has given a mere 9%.

An all-or-nothing theory of right and wrong would have to say that there was a *threshold*, e.g., at 10%, such that if one chose to give 9% one would be wrong, whereas if one chose to give 10% one would be right. If this distinction is to be interesting, it must say that there is a *big* difference between right and wrong, between giving 9% and giving 10%, and a small difference between pairs of right actions, or pairs of wrong actions. The difference between giving 9% and 8% is just the difference between a wrong action and a slightly worse one; and the difference between giving 11% and 12% is just the difference between one supererogatory act and a slightly better one. Given the argument I just rehearsed, the utilitarian should not accept this.⁸

A related reason to reject an all-or-nothing line between right and wrong is that the choice of any point on the scale of possible options as a threshold for rightness will be *arbitrary*. Even maximization is subject to this criticism. One might think that the difference between the best and the next best option constitutes a really significant moral difference, quite apart from the difference in goodness between the options. We do, after all, attach great significance to the difference between winning a race and coming second, even if the two runners are separated by only a fraction of a second. We certainly don't

Reasons without Demands: Rethinking Rightness

attach anything like the same significance to the difference between finishing, say, seventh and eighth, even when a much larger interval separates the runners. True enough, but I don't think that it shows that there really is a greater significance in the difference between first and second than in any other difference. We do, after all, also attach great significance to finishing in the top three. We give medals to the top three and to no others. We could just as easily honor the top three equally and not distinguish between them. When we draw these lines -- between the first and the rest, or between the top three and the rest, or between the final four and the others -- we seem to be laying down arbitrary conventions. And saying that giving 10% is right and giving only 9% is wrong seems analogously conventional and arbitrary.

By contrast, good and bad are scalar concepts, but as with many other scalar concepts, such as rich and tall, we speak of a state of affairs as good or bad (*simpliciter*). This distinction is not arbitrary or conventional. The utilitarian can give a fairly natural account of the distinction between good and bad states of affairs. For example: consider each morally significant being included in the state of affairs. Determine whether her conscious experience is better than no experience. Assign it a positive number if it is, and a negative one if it isn't. Then add together the numbers of all morally significant beings in the state of affairs. If the sum is positive, the state of affairs is good. If it is negative, the state of affairs is bad.

Note that although this gives an account of a real distinction between good and bad, it doesn't give us reason to attach much significance to the distinction. It doesn't make the difference between a minimally good state of affairs and a minimally bad state of affairs more significant than the difference between pairs of good states of affairs or

Reasons without Demands: Rethinking Rightness

between bad states of affairs. To see this, imagine that you are consulted by two highly powerful amoral gods, Bart and Lisa. Bart is trying to decide whether to create a world that is ever so slightly good overall or one that is ever so slightly bad overall. Lisa is trying to decide whether to create a world that is clearly, but not spectacularly, good, or one that is clearly spectacularly good. They each intend to flip a coin, unless you convince them one way or the other in the next five minutes. You can only talk to one of them at a time. It is clearly more important to convince Lisa to opt for the better of her two choices than to convince Bart to opt for the better of his two choices.

However, if utilitarianism only gives an account of goodness, how do we go about determining our moral obligations and duties? It's all very well to know how good my different options are, but this doesn't tell me what morality requires of me. Traditional maximizing versions of utilitarianism, though harsh, are perfectly clear on the question of moral obligation. My obligation is to do the best I can. Even a satisficing version can be clear about how much good it is my duty to produce. How could a utilitarian, or other consequentialist, theory count as a moral theory, if it didn't give an account of duty and obligation? After all, isn't the central task of a moral theory to give an account of moral duty and obligation?

Utilitarians, and consequentialists in general, seem to have agreed with deontologists that their central task was to give an account of moral obligation. They have disagreed, of course, sometimes vehemently, over what actually is morally required. Armed with an account of the good, utilitarians have proceeded to give an account of the right by means of a simple algorithm from the good to the right. In addition to telling us what is good and bad, they have told us that morality requires us to produce a certain

Reasons without Demands: Rethinking Rightness

amount of good, usually as much as possible, that we have a moral obligation to produce a certain amount of good, that any act that produces that much good is right, and any act that produces less good is wrong. And in doing so they have played into the hands of their deontological opponents.

A deontologist, as I said earlier, is typically concerned with such properties of an action as whether it is a killing of an innocent person, or a telling of a lie, or a keeping of a promise. Such properties do not usually come in degrees. (A notable exception is raised by the so-called duty of beneficence.) It is hard, therefore, to construct an argument against particular deontological duties along the lines of my argument against particular utility thresholds. If a utilitarian claims that one has an obligation to produce x amount of utility, it is hard to see how there can be a significant utilitarian distinction between an act that produces x utility and one that produces slightly less. If a deontologist claims that one has an obligation to keep one's promises, a similar problem does not arise. Between an act of promise-keeping and an alternative act that does not involve promise-keeping, there is clearly a significant deontological distinction, no matter how similar in other respects the latter act may be to the former. A utilitarian may, of course, claim that he is concerned not simply with utility, but with maximal utility. Whether an act produces at least as much utility as any alternative is not a matter of degree. But why should a utilitarian be concerned with maximal utility, or any other specific amount?

To be sure, a utilitarian cannot produce an account of duty and obligation to rival the deontologist's, unless he claims that there are morally significant utility thresholds. But why does he want to give a rival account of duty and obligation at all? Why not

Reasons without Demands: Rethinking Rightness

instead regard utilitarianism as a far more radical alternative to deontology, and simply reject the claim that duties or obligations constitute any part of fundamental morality, let alone the central part? My suggestion is that utilitarianism should be treated simply as a theory of the goodness of states of affairs and of the comparative value of actions, which rates alternative possible actions in comparison with each other. This system of evaluation yields information about which alternatives are better than which and by how much. In the example of the doctor, this account will say that the best thing to do is to go and help with the epidemic, but it will say neither that he is required to do so, nor that he is completely unstained morally if he fails to do so.

If a utilitarian has an account of goodness and badness, according to which they are scalar phenomena, why not say something similar about right and wrong: that they are scalar phenomena but that there is a point (perhaps a fuzzy point) at which wrong shades into right? Well, what would that point be? I said earlier that differences in goodness should be reflected by differences in rightness. Perhaps the dividing line between right and wrong is just the dividing line between good and bad. There are two reasons to reject this suggestion. The first is that it seems to collapse the concepts of right and wrong into those of good and bad respectively, and hence, to make the former pair redundant. The second is that, on the account of good and bad states of affairs I offered the utilitarian, it is not clear that there is any satisfactory account of the difference between good and bad *actions* (as opposed to *states of affairs*) with which to equate the difference between right and wrong actions. I do not here have the space to defend this claim, though I have done so extensively elsewhere.⁹

If utilitarianism is interpreted as a scalar theory, that doesn't issue any demands at

Reasons without Demands: Rethinking Rightness

all, it clearly can't be criticized for being too demanding. Does this mean that the scalar utilitarian must agree with the critic who claims (i) we are not frequently required to sacrifice our own interests for the good of others; (ii) we really do have a (fairly) wide range of moral freedom, and (iii) there really are times when we can go above and beyond the call of duty? Strictly speaking, the answers are 'yes' to (i), 'no' to (iii), and 'it depends' to (ii). (i) It may frequently be better to sacrifice our interests for the good of others than to perform any action that preserves our interests. Sometimes it may be much better to do so. However, these facts don't entail any further facts to the effect that we are *required* to do so. (ii) If the claim that we have a wide range of moral freedom is simply the claim that morality doesn't demand only one course of action in most situations, then scalar utilitarians can agree with this. If, on the other hand, moral freedom is supposed to entail not only that morality doesn't narrow down our options with demands, but that we are frequently faced with a wide array of equally choiceworthy alternatives, scalar utilitarians will quite rightly deny this. (iii) As for supererogation, the scalar utilitarian will deny the existence of duty as a fundamental moral category, and so will deny the possibility of actions that go 'beyond' our duty, in the sense of being better than whatever duty demands. The intuition that drives the belief in supererogation can, however, be explained in terms of actions that are considerably better than what would be expected of a reasonably decent person in the circumstances.

At this point, someone might object that I have thrown out the baby with the bath water. To be sure, scalar utilitarianism isn't too demanding, it's not nearly demanding enough! How can a theory that makes no demands fulfill the central function of morality, which is to guide our actions? I turn to this question in the next section.

4. Rightness and Goodness as Guides to Action

Utilitarianism should not be seen as giving an account of right action, in the sense of an action *demand*ed by morality, but only as giving an account of what states of affairs are good and which actions are better than which other possible alternatives and by how much. The fundamental moral fact about an action is how good it is relative to other available alternatives. Once a range of options has been evaluated in terms of goodness, all the morally relevant facts about those options have been discovered. There is no further fact of the form 'x is right', 'x is to-be-done', or 'x is demanded by morality'.

This is not to say that it is a bad thing for people to use the phrases such as 'right', 'wrong', 'ought to be done', or 'demanded by morality', in their moral decision-making, and even to set up systems of punishment and blame which assume that there is a clear and significant line between right and wrong. It may well be that societies that believe in such a line are happier than societies that don't. It might still be useful to employ the notions of rightness and wrongness for the purposes of everyday decision-making. If it is practically desirable that people should think that rightness is an all-or-nothing property, my proposed treatment of utilitarianism suggests an approach to the question of what function to employ to move from the good to the right. In different societies the results of employing different functions may well be different. These different results will themselves be comparable in terms of goodness. And so different functions can be assessed as better or worse depending on the results of employing them.

It is clear that the notions of right and wrong play a central role in the moral thinking of many. It will be instructive to see why this so. There are two main reasons

Reasons without Demands: Rethinking Rightness

for the concentration on rightness as an all-or-nothing property of actions: (i) a diet of examples which present a choice between options which differ greatly in goodness; (ii) the imperatival model of morality. Let's consider (i). When faced with a choice between helping a little old lady across the road, and mugging her, it is usually much better to help her across the road. If these are the only two options presented, it is easy to classify helping the old lady as the 'right' thing to do, and mugging her as 'wrong'. Even when there are other bad options, such as kidnapping her or killing and eating her, the gap between the best of these and helping her across the road is so great that there is no question as to what to do. When we move from considering choices such as these to considering choices between options which are much closer in value, such as helping the old lady or giving blood, it is easy to assume that one choice must be wrong and the other right.

Let us move now to (ii). Morality is commonly thought of as some sort of guide to life. People look to morality to tell them what to do in various circumstances, and so they see it as issuing commands. When they obey these, they do the right thing, and when they disobey, they do a wrong thing. This is the form of some simple versions of divine command ethics and some other forms of deontology. Part of the motivation for accepting such a theory is that it seems to give one a simple, easily applicable practical guide. Problems arise, of course, when someone finds herself in a situation in which she is subject to two different commands, either of which can be obeyed, but not both. In these cases we could say that there is a higher-order command for one rather than the other to be done, or that the agent cannot help doing wrong. The effect of allowing higher-order commands is to complicate the basic commands, so 'do not kill' becomes 'do

Reasons without Demands: Rethinking Rightness

not kill, unless...'. The effect of allowing that there could be situations in which an agent cannot help doing wrong is to admit that morality may not always help to make difficult choices. In either case, one of the motivations for accepting an imperatival model of morality—simplicity, and thus ease of application—is undermined.

Unless one does espouse a simple form of divine command theory, according to which the deity's commands should be obeyed just because they are the deity's commands, it seems that the main justification for the imperatival model of morality is pragmatic. After all, if we don't have the justification that the commands issue from a deity, it is always legitimate to ask what grounds them. That certain states of affairs are good or bad, and therefore should or should not be brought about, seems like a far more plausible candidate to be a fundamental moral fact than that someone should act in a certain way. However, it is generally easier to make choices if one sees oneself as following instructions. It may well be, then, that the imperatival model of morality, with the attendant prominence of the notions of right and wrong, has a part to play at the level of application. It may in fact be highly desirable that most people's moral thinking is conducted in terms of right and wrong. On the other hand, it may be desirable that everyone abandon the notions of right and wrong. I do not wish to argue for either option here, since the issue could probably only be settled by extensive empirical research.

The approach of the last few paragraphs might seem merely to relocate a problem to a different level. I have been claiming that, although morality doesn't actually tell us what we ought to do, there may be pragmatic benefits in adopting moral practices that include demands. Societies that adopt such practices may be better (happier, more flourishing, etc.) than those that don't. But surely this doesn't solve anything. We want

Reasons without Demands: Rethinking Rightness

to know whether we *ought* to adopt such practices. Scalar utilitarianism seems to be silent on that question. Since scalar utilitarianism doesn't tell us what we ought to do, it can't guide our actions (including our choices of what moral practices to adopt and/or encourage in society). But any adequate moral theory must guide our actions. Therefore the theory should be rejected. This argument has three premises:

1. If a theory doesn't guide our action, it is no good.
2. If a theory doesn't tell us what we ought to do, it doesn't guide our action.
3. Utilitarianism, as I have described it, does not tell us what we ought to do.

To assess this argument we need to disambiguate its first premise. The expression 'guide our action' can mean several things. If it means 'tell us what we ought to do' then premise (1) is question-begging. I shall construe it to mean something more like, 'provide us with reasons for acting'. On that reading, I shall concede (1), and shall argue that (2) is false.

Here is Sidgwick in defence of something like (2):

Further, when I speak of the cognition or judgement that 'X ought to be done' -- in the stricter ethical sense of the term ought -- as a 'dictate' or 'precept' of reason to the persons to whom it relates, I imply that in rational beings as such this cognition gives an impulse or motive to action: though in human beings, of course, this is only one motive among others which are liable to conflict with it, and is not always -- perhaps not usually -- a predominant motive. (Sidgwick, 1981, page 34)

As Sidgwick acknowledges, this reason can be overridden by other reasons, but when it is, it still exerts its pull in the form of guilt or uneasiness.

Sidgwick's point rests on internalism, the view that moral beliefs are essentially motivating. Internalism is controversial. Instead of coming down on one side or the

Reasons without Demands: Rethinking Rightness

other of this controversy, I shall argue that, whether one accepts internalism or externalism, the fact that a state of affairs is bad gives reason to avoid producing it as much as would the fact that producing it is wrong.

Suppose internalism is correct. In that case the belief that an act is wrong gives one a reason not to do it. Furthermore, such a reason is necessarily a motivating reason.¹⁰ It seems that the utilitarian internalist should take the position that the belief that a state of affairs is *bad* is also a motivating reason to avoid producing it, and the belief that one state of affairs is *better than the other* may well give the believer a stronger reason to produce the first than the second. If the fact that an act is wrong gives us reason to avoid it, then the fact that it involves the production of a bad state of affairs, by itself, gives us reason to avoid it.

Now let's suppose externalism is true. In that case the fact that an act is wrong gives one a motivating reason to avoid doing it *if one cares about avoiding wrongdoing*. If this is what wrongness amounts to, then it seems no defect in a theory that it lacks a concept of wrongness. For it may be true that one cannot consistently want to avoid doing wrong, believe that an act is wrong and do the act without feeling guilt. But this doesn't provide a distinctive account of wrongness, because we can replace each occurrence of the word 'wrong' and its cognates in the above sentence with other moral terms such as 'an action which produces less than the best possible consequences' or 'much worse than readily available alternatives' and the principle remains true. If the agent cares about doing the best he can, then he will be motivated to do so, feel guilt if he doesn't, and so on. It is true that few of us care about doing the best we can. But then, many of us do not care about doing what we ought either.¹¹

Reasons without Demands: Rethinking Rightness

Whether internalism is correct or not, it looks as if premise (2) in the above argument is false. Abolishing the notion of 'ought' will not seriously undermine the action-guiding nature of morality. The fact that one action is better than another gives us a moral reason to prefer the first to the second. Morality thus guides action in a scalar fashion. This should come as no surprise. Other action-guiding reasons also come in degrees. Prudential reasons certainly seem to function in this way. My judgement that pizza is better for me than cauliflower will guide my action differently depending on how much better I judge pizza to be than cauliflower. Whether moral facts are reasons for all who recognize them (the debate over internalism) is an issue beyond the scope of this paper, but whether they are or not, the significance each of us gives to such moral reasons relative to other reasons, such as prudential and aesthetic reasons, is not something which can be settled by a moral theory.

There are two other reasons I have encountered for requiring utilitarianism to provide an account of the right. The first might be expressed like this: "If utilitarianism is not a theory of the right, it must only be a theory of the good. Likewise, different consequentialist theories will be different theories of the good. But then how do we explain the difference between consequentialist and non-consequentialist theories in general? Since there are no restrictions on the kind of good that any particular version of consequentialism may be a theory of, we are left with nothing that is distinctive about consequentialism."¹²

This is not correct. I can still claim this distinctive feature for consequentialism: it includes the view that the relative value of an action depends entirely on the goodness of its consequences. Of the acts available to the agent, the best action will be the one that

Reasons without Demands: Rethinking Rightness

produces the best consequences, the next best will be the one that produces the next best consequences, and so on. I can also claim that the better the action, the stronger the moral reason to perform it. This is not to concede the point to my opponents. The fact that there is a moral reason to perform some action, even that there is more moral reason to perform it than any other action, doesn't mean that one ought to perform it. (Most of us would acknowledge that one has more moral reason to behave in a supererogatory fashion than simply to do one's duty.) This distinguishes consequentialism from deontology, which allows that one may have a stronger moral reason to perform an action which produces worse consequences. For example, if faced with a choice between killing one and letting five die, the deontologist may acknowledge that five deaths are worse than one, but insist that the better behavior is to allow the five to die. According to that view, morality provides stronger reasons for allowing five deaths than for killing one.¹³ One advantage of the suggestion I offer here over, say, the view that it is of the essence of consequentialism to insist that the agent ought always to do whatever will produce the best consequences, is that it allows satisficing consequentialists and scalar consequentialists to count as consequentialists.

I have also encountered the following reason for requiring utilitarianism to provide an account of the right as well as the good: The utilitarian will have to provide a function from the good to the right in order to compare her theory with various deontological alternatives. Our chief method for comparing moral theories, according to this suggestion, consists in comparing their judgements about which acts are right or wrong. It is true that contemporary discussions of the relative merits of utilitarianism and deontology have often focused on particular examples, asking of the different theories

what options are right or wrong. However, to assume that a moral theory must provide an account of the right in order to be subjected to critical scrutiny begs the question against my proposed treatment of utilitarianism. That utilitarians have felt the need to provide accounts of rightness is testimony to the pervasion of deontological approaches to ethics. Part of what makes utilitarianism such a radical alternative to deontology, in my view, is its claim that right and wrong are not fundamental ethical concepts.

5. Rightness as an Ideal

In this paper, I have argued that utilitarianism is best conceived as a theory of the good, that judges actions to be better or worse than possible alternatives, and thus provides reasons for actions. I have argued that the traditional utilitarian account of rightness as an all-or-nothing property, whether the maximizing or satisficing version, should be abandoned. However, there may be an alternative account of rightness that is particularly congenial to a utilitarian approach. If, instead of conceiving of rightness as a standard that *must* be met (perhaps to avoid censure), we conceive of it as an ideal to which we aspire, we may be able to accommodate it within a scalar framework. The suggestion is that the ideally right action is the maximizing action, and alternatives are more or less right, depending on how close they come to maximizing. Although the ideal itself is often difficult to attain, the theory cannot be charged with being too demanding, since it doesn't include the demand that one attain the ideal. Nonetheless, the ideal functions as a guide. This would be similar to the approach taken by many Christians, who view Christ as a moral exemplar. A common articulation of this view is the question 'what would Jesus do?', often abbreviated on bracelets, bumper stickers, handguns, and the like as

Reasons without Demands: Rethinking Rightness

‘WWJD’. Inasmuch as the extant accounts of Christ’s life provide a basis for answering this question, the answer is clearly supposed to function as an ideal towards which we are supposed to aspire, and not as a demand that must be met in order to avoid wrongdoing. The closer we come to emulating the life or the actions of Christ, the better our lives or our actions are.

The utilitarian version (WWJSMD?) might be easier to apply, both epistemically and practically. There are, of course, well-known epistemic problems with even a subjective expected-utility version of utilitarianism, but these pale into insignificance compared with the difficulty of figuring out what Jesus would do, whether the (presumably) actual historical figure, or the literary composite portrayed in the biblical (and other) sources. As for the practical problems with viewing Christ as an exemplar, it may turn out that the ideal is not simply difficult to attain, but in some cases impossible. On the assumption that Christ had divine powers, an assumption that is undoubtedly accepted by most adherents of the Christ-as-exemplar moral theory, we may sometimes be literally unable to do what Jesus would have done. For example, suppose I am attending a wedding in Lubbock (TX), and the wine runs out. Amid the wailing and the gnashing of teeth I glance at the ‘WWJD’ engraved on my cowboy boots. Well, it’s clear what Jesus would do in this case (John: 2, 1-10), but I simply can’t turn water into wine. However, the utilitarian ideal is, by definition, possible. In this case it might involve driving outside the city limits (Lubbock is dry in more than one sense) to one of the drive-through liquor stores, loading up on the surprisingly good local wines, and returning to spread cheer and much-needed intoxication to the wedding festivities. Or perhaps, more plausibly, it might involve sending the money to famine relief.

6. Prohibitions and the ‘Separateness of Persons’

For those who are inclined to think that traditional maximizing utilitarianism is seriously threatened by the objection that it is too demanding, the suggestion that we interpret the theory in scalar fashion, either abandoning the notion of rightness altogether or interpreting it as an ideal, may be particularly attractive. There are also, as I have argued, independent reasons for adopting a scalar version of utilitarianism.¹⁴ However, adopting a scalar version of utilitarianism does not, as far as I can see, have any bearing on Peter Vallentyne’s second line of criticism of consequentialist theories, that they don’t include constraints required to recognize the ‘normative separateness of persons’. I will, therefore, close with a brief discussion of why I don’t think a utilitarian (any version) should be worried by this line of attack on the theory.

The criticism that utilitarianism does not recognize or account for the ‘separateness of persons’ has become commonplace since Rawls, but what exactly does it mean? Peter Vallentyne’s explanation is fairly representative:

[I]ndividuals have certain rights that may not be infringed simply because the consequences are better. Unlike prudential rationality, morality involves many distinct centers of will (choice) or interests, and these cannot simply be lumped together and traded off against each other.

The basic problem with standard versions of core consequentialism is that they fail to recognize adequately the *normative separateness of persons*. Psychological autonomous beings (as well, perhaps, as other beings with moral standing) are not merely means for the promotion of value. They must be *respected* and *honored*, and

Reasons without Demands: Rethinking Rightness

this means that at least sometimes certain things may not be done to them, even though this promotes value overall. An innocent person may not be killed against her will, for example, in order to make a million happy people significantly happier. This would be sacrificing her for the benefit of others.

There seem to be several distinct ideas here. (i) Individuals have rights, that at least sometimes trump utility calculations. (ii) Individuals' interests can't simply be traded off against each other. (iii) Individuals must be *respected* or *honored*. Consider these claims in reverse order. A utilitarian may claim, with some justification, that the demand for equal consideration of interests embodied in her theory (and other consequentialist theories) is precisely what it means to respect or honor individuals. It is only when I weigh your interests equally with the interests of all others whom I can affect that I adequately respect or honor you. Deontological constraints function to disallow the consideration of certain interests in certain circumstances. Thus they, at least sometimes, prevent us from respecting or honoring certain individuals.

At this point, the critic of utilitarianism will no doubt claim that I have (perhaps willfully) misunderstood (iii). In fact, he might claim that (i) and (ii) explain *what it means* to honor or respect individuals. (ii) denies the aggregative feature of utilitarianism. The problem with the denial of tradeoffs or aggregation is that even committed anti-consequentialists accept them in many circumstances. For example, suppose that Homer is faced with the painful choice between saving Barney from a burning building, or saving both Moe and Apu from the building. Clearly it is better for Homer to save the larger number, *precisely because it is a larger number*. The proponent of (ii) might try to accommodate this intuition by limiting the scope of tradeoffs. For

example, perhaps we are allowed to trade lives for lives (or similarly serious harms), but we are not allowed to trade lives for convenience.¹⁵ Homer can save the lives of Moe and Apu rather than Barney, but he can't leave Barney to die in order to provide all the inhabitants of Springfield with a few minutes extra free time every day.¹⁶ However, any such attempt to limit the scope for tradeoffs faces at least two serious problems. First, such a move almost certainly entails denying the transitivity of 'all-things-considered better than'. Second, we commonly accept tradeoffs between lives and much lesser values, such as convenience. For example, we allow public projects such as building a bridge in order to make travel between two places more convenient, even when we know that several people will die in the course of the construction. Likewise, even most anti-consequentialists don't demand that highway speed limits be lowered to the optimal point for saving lives, even though the advantages of higher speed limits are increased convenience for many.¹⁷

We are left with (i), the claim that individuals have rights that sometimes trump utilities. Utilitarianism is criticized for failing to distinguish between the following pair of cases (adapted from Foot (1984)): (a) Homer must choose whether to save Barney, who is trapped on one side of Springfield, or both Moe and Apu, who are trapped on the other side. He can't save all three, and no-one else can save any of them. (b) Homer, and no-one else, can save both Moe and Apu, who are trapped on the edge of Springfield Gorge. However, in order to reach them in time to save them, he must run over and kill Barney, who is trapped on a narrow segment of the only road leading to the gorge. We are supposed to agree that Homer may choose to save Moe and Apu in (a), but not in (b). If he saves Moe and Apu in (b), he will violate Barney's right not to be killed. But don't

Reasons without Demands: Rethinking Rightness

Moe and Apu have the right to be saved? Perhaps, but, if so, it is not as important (strict, stringent, etc.) as Barney's right not to be killed. In general, if the rights view is to present a genuine alternative to consequentialism, negative rights not to be harmed in some way must be stronger than the corresponding positive rights, if any, to be aided in avoiding such harm. More specifically, the duty not to harm in a certain way must be stricter than the corresponding duty to prevent such harm. Claims that negative rights and duties are (at least usually) stronger than positive rights and duties will have to be grounded in an account of the alleged moral significance of the general distinction between doing and allowing, of which the distinction between killing and letting die is a specific example. This topic is the subject of much debate, which I don't have the space here to recapitulate.¹⁸ It is, however, no surprise that consequentialists deny the moral significance of the doing/allowing distinction.¹⁹ If, as I have only briefly suggested here, the criticism that consequentialism does not recognize the 'normative separateness of persons' really amounts to the claim that consequentialism does not endow the doing/allowing distinction with intrinsic moral significance, no consequentialist should be troubled by it.

ALASTAIR NORCROSS

REFERENCES

Bennett, J. (1995). *The act itself*. Oxford: Oxford University Press.

Feinberg, J. (1961) "Supererogation and Rules," *Ethics* 71, 276-288.

Foot, P. "Killing and Letting Die" (1984), repr. in Steinbock and Norcross, *Killing and*

Letting Die, 280-289.

Howard-Snyder, F. and Norcross, A. (1993) "A Consequentialist Case for Rejecting the Right", *The Journal of Philosophical Research*, 18, 109-125.

Norcross, A. (1995) "Should Utilitarianism Accommodate Moral Dilemmas?", *Philosophical Studies*, 79:1, 59-85.

Norcross, A. (1997a) "Good and Bad Actions", *The Philosophical Review*, 106:1, 1-34.

Norcross, A. (1997b) "Comparing Harms: Headaches and Human Lives", *Philosophy and Public Affairs*, 26:2, 135-167.

Norcross, A. (2002) "Contractualism and Aggregation", *Social Theory and Practice*, 28:2, 303-314.

Norcross, A. (2003) "Killing and Letting Die", in R.G. Frey and C. H. Wellman (Eds). *The Blackwell Companion to Applied Ethics*, Blackwell.

Norcross, A. (2004), "Scalar Act-Utilitarianism", in H. West (ed). *Blackwell Guide to Mill's Utilitarianism*, Blackwell.

Rawls, J. (1971) *A Theory of Justice*, Cambridge, Mass.: Belknap Press of Harvard University Press.

Scanlon, T.M. (1998). *What we owe to each other*. Cambridge, Mass.: Belknap Press of Harvard University Press.

Sidgwick, H. (1981) *The Methods of Ethics*, 7th. Ed. Hackett.

Slote, M. (1985) *Common-sense Morality and Consequentialism* Boston: Routledge and Kegan Paul.

Steinbock, B. and Norcross, A. (eds). (1994), *Killing and Letting Die*, 2nd edn. New York: Fordham University Press.

SUGGESTIONS FOR FURTHER READING

Howard-Snyder, F. (1994) “The Heart of Consequentialism”, *Philosophical Studies*, 76:1, 107-129.

Kagan, S. (1989). *The limits of morality*. New York: Oxford.

Kagan, S. (1998). *Normative ethics*. Boulder: Westview Press.

Parfit, D. (1984). *Reasons and Persons*. Oxford: Oxford University Press.

Scheffler, S. (ed.) (1988). *Consequentialism and its critics*. Oxford: Oxford University Press.

Smart, J.J.C. and Williams B. (1973) *Utilitarianism: For and Against* Cambridge: Cambridge University Press.

Unger, P. (1996) *Living High and Letting Die: Our Illusion of Innocence*, Oxford University Press.

¹ Given certain ways of stating the second and third objections, my theory doesn't answer them. Peter Vallentyne's statements of these criticisms of utilitarianism are “(2) it leaves agents inadequate moral freedom (judges too few actions morally permissible), and (3) it leaves no room for permissible actions that are morally better than other permitted actions.” On my approach, *no* actions are permissible. However, since no actions are impermissible either, the spirit of the criticisms clearly doesn't apply.

² I am concerned only with theories which are agent-neutral, and whose value theories are relatively fine-grained.

³ My apologies to the proponents of virtue ethics, the third-party candidate of ethical theories.

⁴ See, for example, Feinberg (1961).

⁵ Slote (1985), chapter 3, discusses this suggestion.

⁶ Though the approach of W. D. Ross might plausibly be interpreted in a scalar fashion.

⁷ I take the term 'scalar' from Slote (1985), who discusses scalar morality in chapter 5.

⁸ It might be objected that maximizing utilitarianism does in fact give a scalar account of wrongness, if not of rightness. Some actions are closer to being right than are others, and so are less wrong. However, the claim that an action is closer to the best action than is another is quite consistent with the claim that it is no less wrong than the latter.

⁹ Norcross (1997a).

¹⁰ There can be reasons that are not necessarily motivating, e.g. prudential reasons. You may have a prudential reason to act in a certain way, be aware of the reason, and yet be not in the least motivated so to act. I am not here thinking of cases in which other motivations--moral, aesthetic, self-indulgent and the like--simply overwhelm prudential motivations. In such cases you would still be motivated to act prudentially, but more motivated to act in other ways. If you simply didn't care about your own well-being, prudential reasons would not be in the least motivating. But someone who didn't care about her own well-being could still have, and even be aware of, prudential reasons. Similarly, if you are asked what is the sum of five and seven, you have a reason to reply 'twelve', but you may be not in the least motivated to do so, for you may not care about arithmetic truth, or any other truth. There may be reasons other than moral reasons that are necessarily motivating. For example, the belief that a particular action is the best way to satisfy one of your desires may provide a necessarily motivating reason to perform that action. The motivation may be outweighed by other motivations.

¹¹ Slote (1985) points this out.

¹² I have heard this objection from Daniel Howard-Snyder and Shelly Kagan.

¹³ The full story about what distinguishes consequentialism from deontology will have to be more complicated than this. It may, for example, incorporate the claim that the consequentialist ranking of states of affairs is not agent-centered. See Scheffler (1982) for a discussion of this notion. On the other hand, we may wish to maintain (as Peter Vallentyne does in his contribution to this volume) that an agent-relative value theory may be incorporated into a consequentialist structure to give a form of consequentialism. Whether we classify, for example, ethical egoism as a form of consequentialism or as an entirely different form of moral theory (or not as a moral theory at all) seems to me to be of very little interest.

¹⁴ For more discussion of these and other reasons see, for example, Howard-Snyder and Norcross (1993), Norcross (1997a), Norcross (2004).

¹⁵ Scanlon tries such a move in his 1998. I critique it in Norcross (2002).

¹⁶ I leave the reader to fill in the details of this and other examples involving the endlessly fascinating inhabitants of Springfield.

¹⁷ For detailed discussion of both these points, see Norcross (1997b).

¹⁸ See, for example, Bennett (1995), Steinbock and Norcross (1994), Norcross (2003).

¹⁹ It may be possible to construct a consequentialist theory that is sensitive to this distinction (see Norcross (1995)), but I know of no-one who embraces such a theory.