Chapter 10

The Scalar Approach to Utilitarianism

Alastair Norcross

Introduction

Consequentialist theories such as utilitarianism have traditionally been viewed as theories of right action. Consequentialists 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 dominant consequentialist function from the good to the right, at least since Sidgwick, has been maximization: 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. Consequentialists 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 deontology. But this is not the case with consequentialism. Goodness and badness, especially in the utilitarian value theory, 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?

I will argue that, from the point of view of a consequentialist, actions should be evaluated purely in terms that admit of degrees. My argument is directed towards those who are already attracted to consequentialism. I am not undertaking to argue for consequentialism against any rival moral theory. Perhaps, though, some who have rejected consequentialism because of its views on rightness and wrongness may be more attracted to a theory constructed along the lines that I suggest.

I shall conduct my discussion in terms of utilitarianism, since this is the most popular form of consequentialism, and the one to which I adhere. However, since none of my points rely on a specifically utilitarian value theory, my argument applies quite generally to consequentialist theories.

The Demandingness Objection

Since, according to maximizing utilitarianism, any act that fails to maximize is wrong, there appears to be no place for actions that are morally admirable but not required, and agents will often be required to perform acts of great self-sacrifice. This gives rise to the common charge that maximizing 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 interests for the good of others, at least not as much as utilitarianism tells us. Furthermore, 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 non-human 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 non-human animals merely on its conflict with "common sense." Furthermore, it is worth noting that one of the main contemporary alternatives to a (roughly) consequentialist approach to morality is often referred to as "common-sense morality." 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. 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. For example, a doctor, who hears of an epidemic in a remote village 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 (Feinberg 1961: 276–88). Such an action does not seem to be 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 (Slote 1985: ch. 3). 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, 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, at least as far as right and wrong are concerned.

Scalar Utilitarianism

Here's an argument for the view that rightness and wrongness isn't an all-or-nothing affair. Let us call this, following Slote (1985: ch. 5) "scalar morality." 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 percent of his income to charity. The difference between giving 8 percent and 9 percent is approximately the same, in some obvious physical sense, as the difference between giving 9 percent and 10 percent, or between giving 11 percent and 12 percent. Such similarities should be reflected in moral similarities. A moral theory which says that there is a *really significant* moral difference between giving 9 percent and 10 percent, but *not* between giving 11 percent and 12 percent, looks misguided. At least, no utilitarian should accept this. She will be equally concerned about the difference between giving 11 percent and 12 percent as the difference between giving 9 percent and 10 percent. To see this, suppose that Jones were torn between giving 11 percent and 12 percent and that Smith were torn between giving 9 percent and 10 percent. 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 percent) would satisfy this goal as well as an extra $5,000 from Smith (who has given 9 percent). It does not matter whether the $5,000 comes from one who has already given 11 percent or from one who has given a mere 9 percent.

An all-or-nothing theory of right and wrong would have to say that there was a *threshold*, e.g., at 10 percent, such that if one chose to give 9 percent one would be wrong, whereas if one chose to give 10 percent 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 percent and giving 10 percent, and a small(er) difference between similarly spaced pairs of right actions, or pairs of wrong actions. The difference between giving 9 percent and 8 percent is just the difference between a wrong action and a slightly worse one; and the difference between giving 11 percent and 12 percent is just the difference between one supererogatory act and a slightly better one. In fact, if the difference between right and wrong is at all significant, it must be possible for it to offset at least some differences in goodness. For example, if the threshold for rightness in a case of charitable giving

were $10,000 for both Smith and Jones, the difference between Smith giving $10,000 and giving $9,000 must be more significant than the difference between Jones giving $9,000 and giving somewhat less than $8,000. But suppose that Smith is wavering between giving $10,000 and giving $9,000, and that Jones is wavering between giving $9,000 and giving $7,999. No utilitarian would consider it more important to persuade Smith to give the higher amount than to persuade Jones to give the higher amount. This doesn't just apply to actions. No utilitarian would consider it better that (or would hope more fervently that, etc.) Smith opts for the higher amount than that Jones opts for the higher amount.

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 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 be laying down arbitrary conventions. And saying that giving 10 percent is right and giving only 9 percent 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 or worse than no experience. Assign it a positive number if it is better, and a negative one if it is worse. Then add together the numbers of all morally significant beings in the state of affairs. If the sum if 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 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 (+1) or one that is ever so slightly bad overall (–1). Lisa is trying to decide whether to create a world that is clearly, but not spectacularly, good (+1), or one that is clearly *spectacularly* good (+1 *billion*). 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. Isn't it 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, at least in the twentieth century onwards, 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 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 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. In fact, Mill's statement of the principle of utility might seem to suggest such an approach:

> The creed which accepts as the foundation of morals, Utility, or the Greatest Happiness Principle, holds that actions are right in proportion as they tend to promote happiness, wrong as they tend to produce the reverse of happiness. (Mill 1861, reprinted as Part II of this volume, Ch. II, para. 2; citations are simply by roman numeral for chapter and arabic numeral for paragraph)

There are, of course, notorious difficulties with interpreting Mill's statement. Talk of tendencies of actions to promote happiness, for example, could lead to a form of rule-utilitarianism. However, since my concern in this paper is not to offer an interpretation of Mill, I will set aside issues of scholarship, and focus instead on the suggestion that an action is right just in case it is good, and wrong just in case it is 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 with which to equate the difference between right and wrong actions. (See Norcross 1997.)

Wrongness as Blameworthiness

We tend to blame and punish agents for their wrong actions and not for actions that are not wrong. And we tend to consider it wrong to blame or punish someone for an action that was not wrong. Perhaps, then, a utilitarian (or other theorist) can analyze wrongness in terms of blameworthiness. Consider this well-known passage from Mill:

> We do not call anything wrong, unless we mean to imply that a person ought to be punished in some way or other for doing it; if not by law, by the opinion of his fellow-creatures; if not by opinion, by the reproaches of his own conscience. (V, 14)

Although, as I said earlier, my main concern is not Mill scholarship, it is worth noting that a close reading of the paragraphs immediately preceding this passage suggests that Mill is not here proposing an *analysis* of wrongness that fits with utilitarianism. He is, rather, pointing out some features of the ordinary usage of the term "wrong." Nonetheless, an analysis of wrongness in terms of punishment is worth considering. This suggests the possibility of a scalar conception of wrongness. Since censure (and other forms of punishment) comes in degrees, then perhaps wrongness might also come in degrees.

Consider the following definition of wrong action:

WA: An action is wrong if and only if it is appropriate to impose various sanctions on the agent.

What does it mean to say that it is "appropriate" to sanction? Since appropriateness is a normative notion, the most natural understanding is to think of it as meaning "obligatory." (Alternatively, we could understand it as meaning "permissible." I shall focus on the obligatory reading here, since it is the most popular. The chief objection to the obligatory reading also applies to the permissibility reading.) In that case, WA would be:

WA1: An action is wrong if and only if we ought to impose various sanctions on the agent.

But if WA is to be understood as WA1, it leads to a definitional circle or regress. It tells us to understand what is wrong in terms of what it is wrong not to do. (I take it that "wrong" and "ought not to be done" are interchangeable.) But we don't have a better grasp on the notion of "ought to sanction" than we have on the notions of "ought to keep promises" or "ought to feed the hungry." Trying to understand the wrongness of one action in terms of the wrongness of other actions is unenlightening.

There is an alternative account of appropriateness according to which it is still normative. That account says that an action is appropriate if and only if it is optimific. WA would then amount to:

WA2: An action is wrong if and only if it is optimific to punish the agent.

This suggestion avoids the uninformative circularity of WA1. Let us suppose that WA2 expresses the sort of connection between wrongness and censure that people have in mind. Can it provide the utilitarian with a distinction between badness and wrongness?

I believe that WA2 is not available to the utilitarian as a way of distinguishing wrong actions from bad actions. For he cannot identify wrong actions with actions which it is optimific to sanction. To see this consider what Sidgwick says about praise.

> From a Utilitarian point of view, as has been before said, we must mean by calling a quality, "deserving of praise," that it is expedient to praise it, with a view to its future production: accordingly, in distributing our praise of human qualities, on utilitarian principles, we have to consider primarily not the usefulness of the quality, but the usefulness of the praise. (Sidgwick 1962: 428)

The utilitarian will, of course, say the same about censure as Sidgwick says about praise: we should assess whether it is good to punish or blame someone by assessing the utility of doing so. Punishing and blaming are actions just like promise-keeping and killing and, like those actions, their value is determined by their consequences, their power to produce utility.

If there is a conceptual connection of the sort asserted by WA2 between an action's being wrong and its being appropriate to punish the agent, and if Sidgwick's account of when it is appropriate to punish is correct, then we should be able to determine whether an act is wrong by determining whether punishing the agent of that act will produce more utility than not. On the other hand, if there is good reason to reject this method of deciding whether someone has done wrong, then there is reason to reject WA2.

I submit that there is reason to reject the claim that the *wrongness* of an action is determined by whether punishing the agent would produce more utility than not. This is because our concept of wrongness is constrained by one or both of the following principles which conflict with WA2:

1: If action *x* is wrong, then an action *y* done by someone in exactly similar circumstances, with the same intention and the same consequences, is also wrong.

We might call this the principle of universalizability. It might, however, be optimific to punish the agent of *x* but not the agent of *y*. Hence, according to WA2, *x* would be wrong, but *y* would not.

2: If someone does the best she can, and does very well indeed, then she has done nothing wrong.

But it can sometimes be optimific to punish a utility-maximizer. For example, imagine that Agnes has always produced as much utility as it was possible for her to produce. Moreover, none of her actions has led to any unfortunate consequences, such as someone's untimely death or suffering. Punishing her as a scapegoat might nevertheless produce more utility than not doing so. It is absurd to say that she has done something wrong just in virtue of the fact that it is appropriate or optimific to punish her.

Given any one of these constraints on any recognizable understanding of wrongness, the utilitarian cannot say that whether an action is wrong is determined by whether it is optimific to punish the agent for doing it. (The principles I have expressed here would produce conflicts with a satisficing version of WA2 as well.)

It might be objected that the conflict between principles (1) and (2) and WA2 can be explained away if we see the relationship between the concepts of wrongness and censure as less simple than the one we have suggested. Instead of saying that wrongness is a necessary and sufficient condition for the appropriateness of censure, it might be urged, we need a connection that allows for exceptions. Why not say that wrongness makes it prima facie appropriate to censure the agent? This prima facie appropriateness can be outweighed by such factors as the agent's motivational state or very extreme consequences of, say, not punishing an innocent person. We recognize intuitively that the appropriateness of punishing or blaming someone for doing wrong is occasionally outweighed by other considerations, e.g., if the agent has suffered a great deal as a result of his action. So the suggestion is:

WA3: An action is wrong if and only if it is prima facie appropriate to blame its agent.

Presumably also, the fact that an agent has done no wrong is a prima facie reason not to punish the agent. Both kinds of reasons can be overridden, particularly in the sorts of bizarre cases that philosophers come up with. This would account for our disinclination to say that Agnes did something wrong, even if we grant that it was appropriate to punish her.

This suggestion will not work, however. The utilitarian does not think it is prima facie appropriate to blame the agent of every wrong action, or prima facie appropriate not to blame someone who has done no wrong. This is because the *sole* determinant of the appropriateness of blame or punishment is how much utility will be produced by doing so. Imagine someone (say, Agnes) who has done no wrong. Suppose that punishing her will produce very slightly more utility than not doing so. The utilitarian will judge that punishing her is better than not. If her innocence did create a presumption against punishing her, then that presumption should be enough to outweigh the very slight gain in utility. But this is not how the utilitarian sees things.

I submit that there can be no conceptual connection, for the utilitarian, between wrongness and punishability or blameworthiness.

Even if we reject a conceptual connection between wrongness and punishment, isn't there a simple conceptual analysis of *rightness* in terms of reasons? If a utilitarian accepts that one possible action is better than another just in case it has better consequences, must she not also accept that there is more (moral) reason to perform the better action? Furthermore, must she not accept that there is the most (moral) reason to perform the *best* action available to the agent? But doesn't "(morally) right" simply mean "supported by the strongest (moral) reasons" or "what we have most (moral) reason to do"? In which case, we seem to have arrived back at the maximizing conception of rightness.

The problem with this suggestion is that it is highly implausible that there is such a simple conceptual connection between rightness and maximal reason. Consider the example of supererogation discussed above. Supererogatory acts are supposed to be morally superior to their merely permissible alternatives. There will be more (moral) reason to perform putative supererogatory acts than to perform their morally inferior alternatives. But if "right" simply means "supported by the strongest reasons," supererogatory acts will be right (unless there are even better alternatives) and the morally inferior, but supposedly acceptable, alternative acts will be wrong. But if "right" and "duty" are equated, then, by definition, there can be no acts that go "above and beyond" duty, in the sense of being better than duty. So it would seem that those who argue for the category of supererogatory actions are simply confused about the meanings of words. But this is highly implausible. Likewise, those who criticize utilitarianism's maximizing account of rightness as being too demanding are surely not simply mistaken as to the meaning of the word "right."

Rightness and Goodness as Guides to Action

If utilitarianism is interpreted as a scalar theory, which doesn't issue any demands at 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 (1) we are not frequently required to sacrifice our own interests for the good of others; and (2) there really are times when we can go above and beyond the call of duty? Strictly speaking, the answers are "yes" to (1), and "no" to (2). (1) 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. (2) 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 bathwater. 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?

Utilitarianism should not be seen as giving an account of right action, in the sense of an action *demanded* 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 for the concentration on rightness as an all-or-nothing property of actions: (1) a long list of examples which present a choice between options which differ greatly in goodness; (2) the imperatival model of morality. Let's consider (1). 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 (2). 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 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 us 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 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 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 defense of something like (2):

> Further, when I speak of the cognition or judgment 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. (1962: 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 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. 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.

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 judgment that cauliflower is better for me than pizza will guide my action differently depending on how much better I judge cauliflower to be than pizza. 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. 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?

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 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 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. 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 judgments 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.

References and further reading

Bennett, Jonathan (1993) "Negation and abstention: two theories of allowing," *Ethics* 104, pp. 75–96.

Feinberg, Joel (1961) "Supererogation and rules," *Ethics* 71, pp. 276–88.

Mill, John Stuart (1861) *Utilitarianism*. Reprinted as Part II of this volume.

Norcross, Alastair (1997) "Good and bad actions," *Philosophical Review* 106, pp. 1–34.

Sidgwick, Henry (1962) *The Methods of Ethics*, 7th edn. Chicago: University of Chicago Press.

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