

CHAPTER 6

Acts or rules? The fine-tuning of utilitarianism

Brad Hooker

Peter Singer's contributions to moral philosophy are difficult to exaggerate. In particular, his discussions of moral duties concerning poverty, animals, abortion, and euthanasia are immensely influential. The books he has authored are among the best sellers of contemporary philosophy, and even the list of those he has edited is strikingly impressive. I believe that both contemporary moral practice and contemporary moral theory are very much better than they would be without him.

What are the qualities that have made his published work so successful? Singer's writing is wonderfully direct, engaging, and clear. He addresses the relevant complexities but avoids any inessential tangents. And he puts forward extremely powerful arguments. Even where people cannot bring themselves to accept Singer's conclusions, Singer's arguments *must* be addressed.

Singer's own moral theory is a form of utilitarianism. Utilitarianism is an approach to morality and social policy that grounds justification in the promotion of utility. Some self-described utilitarians have used the term 'utility' to mean value. Most utilitarians, however, mean something narrower than value. Specifically, utility is meant as aggregate benefit-minus-harm, which is often referred to as *net benefit*, or *welfare*, *well-being*, or *aggregate personal good*.

One of the main attractions of utilitarianism has always been its impartiality. In the calculation of aggregate welfare, each individual's welfare counts equally. Such impartiality, according to utilitarians, involves counting benefits or harms to any individual the same as benefits or harms, of the same size, to anyone else. Singer often calls this kind of impartiality 'equal consideration of interests'.

Although not everyone assumes moral assessment and justification must be fundamentally impartial, the assumption is widespread. It is also understandable. After all, when we try to justify their behaviour to one another, we often try to show that our behaviour complies with moral principles and

considerations. But how could our appeal to moral principles and considerations provide justification to others for our behaviour unless morality has only impartial principles at the foundational level?

Since utilitarianism grounds justification in aggregate welfare, we need to consider the nature of welfare. There are different views about this – hedonistic views, desire-fulfilment or preference-satisfaction views, and perfectionist, eudaimonistic, or objective list views.¹ I will not wade into the debate between these rival views here. But let me remark in passing that I think the debates between advocates of the main different theories of welfare, though fascinating theoretically, are often irrelevant practically, at least in cases where only *human* welfare is at stake. For very often a person's happiness, desire-fulfilment, and eudaimonistic flourishing will be served by the same events.

The distinction between different theories of welfare is less important than the distinction between different utilitarian theories about how moral requirements are related to welfare. The two most important kinds of utilitarian theories about how moral requirements are related to welfare are act-utilitarianism and rule-utilitarianism. Actually, there are many versions of each of these kinds. But I shall focus on the most prominent ones. The most prominent version of act-utilitarianism requires one to do acts that produce at least as much utility as any other acts one could have done instead. The most prominent version of rule-utilitarianism favours the rules whose widespread internalization would maximize utility and then takes acts to be morally permissible, morally required, or morally wrong depending on whether they are permitted, required, or forbidden by these rules.

Both act-utilitarianism and rule-utilitarianism can be formulated in terms of actual utility or in terms of expected utility. Different possible consequences are more or less probable, and more or less beneficial or harmful. To calculate the expected utility of an act or rule, multiply the net benefit or harm of each possible consequence of that act or rule by the probability of that consequence occurring. The numbers produced by these multiplications are then summed in order to generate the expected utility of the act or rule.

As an example, consider the act of Jill's explaining expected utility to Jack. There are many possible consequences of this act. But let us pretend there are only two, Jack's being enlightened or Jack's being befuddled.

¹ The best discussion of the differences among these theories of welfare remains Appendix I of Derek Parfit's *Reasons and Persons* (Oxford: Oxford University Press, 1984).

Suppose the benefit of Jack's being enlightened is +5 and the harm of Jack's being befuddled is -3 and the probability of Jack's being enlightened by Jill's explanation is 60 per cent and the probability of Jack's instead being befuddled by Jill's explanation is 40 per cent. Then the expected utility of Jill's explaining expected utility to Jack is $(+5 \times 0.6) + (-3 \times 0.4)$, which equals $3 + (-1.2)$, which equals 1.8. That was an example concerning acts, but the same sort of calculations can be done to determine the expected utility of this or that possible moral rule.

The calculation of expected utility I have just used as an illustration presumed not only precise quantifications of possible benefits and harms but also precise probability percentages. Such precision, however, is utterly unrealistic. Normally, even the very best available calculation of expected utility is fairly rough, quite speculative, and somewhat impressionistic. To pretend otherwise is ridiculous.

Having explained expected utility, I now state act-utilitarianism and rule-utilitarianism as theories that determine moral wrongness in terms of expected utility. Act-utilitarianism holds that an act is morally required if and because the individual act's expected utility is greater than the expected utility of any alternative individual act. Rule-utilitarianism holds that an act is morally required if and because the act is required by a code of rules whose internalization has greater expected utility than the internalization of any alternative code.

When people first hear about act-utilitarianism, they often assume that act-utilitarianism holds both that

- (a) what makes an act morally required is that the act would *maximize* net expected aggregate utility
- and that
- (b) everyone should normally and routinely *decide how to behave* by calculating the utilities of alternative actions and then choosing the one with the greatest net expected aggregate utility.

Indeed, act-utilitarians do hold what makes an act morally required is that it would *maximize* net expected aggregate utility. But they do not hold that everyone should normally and routinely *decide how to behave* by calculating the utilities of alternative actions and then choosing the one with the greatest net expected aggregate utility. For utility would *not* be maximized by our regularly deciding how to behave by calculating the utilities of alternative actions.

There are at least five reasons why utility would *not* be maximized by our regularly deciding how to behave by calculating the utilities of alternative actions:

- (1) Often we don't know and can't easily find out what benefits or harms would result from an act, or what the probabilities are.
- (2) Even if we can find out, the costs of finding out or of then calculating the utility might well outweigh the extra benefits resulting from the superior act.
- (3) Furthermore, there is the serious risk of making mistakes in the calculation.
- (4) And people's expecting other people routinely to *decide how to behave* by trying to calculate the utilities of alternative actions would cause widespread insecurity and distrust, since people would expect that others would miscalculate because of incompetence or deep-seated bias. For example, if you thought I made my every decision by calculating the utilities, you might worry that I might become convinced that killing you, or stealing from you, or breaking my promises to you, or lying to you would maximize utility.
- (5) In order for people to become reliably disposed to maximize impartial utility in their every decision, they would have to be stripped of all traces of natural partiality towards friends, lovers, and family. But this would effectively eliminate human friendship and make humans 'lousy lovers' (to use C.L. Ten's phrase). If all became lousy lovers and friendless, human happiness would greatly diminish. In Ronald Dworkin's words, 'If we felt nothing more for friends or lovers or colleagues than the most intense concern we could possibly feel for all fellow citizens, this would mean the extinction not the universality of love.'² In other words, humans are capable of very strong concern for some or equal concern for all, but not strong concern for all. So the price of equal concern for all is, to use Sidgwick's quotation from Aristotle, 'but a watery kindness'.³

For these reasons, what would actually maximize expected utility would not be people's being regularly disposed to decide how to behave by calculating the utilities of alternative actions. While act-utilitarianism does indeed hold that what makes required acts required is that they maximize net expected aggregate utility, act-utilitarianism also holds that the required procedure for making everyday moral decisions (especially ones in the heat of the moment, when information might be short and reasoning hasty and biased) is to resort to tried and tested moral rules, such as 'don't hurt others', 'don't

² Ronald Dworkin, *Law's Empire* (Cambridge, MA: Harvard University Press, 1986), 215.

³ Henry Sidgwick, *The Methods of Ethics*, 7th edn (London: Macmillan, 1907), 434.

steal', 'don't break promises', 'don't lie', 'do attend to the needs of your near and dear', etc.

Turn for a moment from the question 'how should individual agents make everyday decisions?' to the question 'what moral rules should people accept?' I mean this to be the question 'what guidelines to decision-making should people accept, given that these guidelines are going to be reflected in (and reinforced by) dispositions to feel guilt when one breaks these rules and to feel resentment or indignation towards others who break them?' Well, here act-utilitarianism and rule-utilitarianism agree that the best rules are those whose internalization would produce the best consequences.

Admittedly, act-utilitarians and rule-utilitarians do not agree entirely about the criteria by which rules should be evaluated. Strictly speaking, act-utilitarianism favours the decision procedure, or rules, the internalization of which by the *individual* agent would maximize expected utility. So, according to act-utilitarianism, since Jack's and Jill's capacities and situations may be very different, the best rules and decision procedures for Jack to accept may be different from the best rules and decision procedures for Jill to accept. Act-utilitarianism tailors rules and decision procedures to different people. That provision goes against the spirit of rule-utilitarianism, a theory committed to the idea that a rule cannot be appropriate unless the rule's being accepted widely (if not universally) would maximize expected utility.

When act-utilitarians endorse the above optimific decision procedure, they acknowledge that following this decision procedure does not guarantee that we will do the act with the best consequences. Sometimes, for example, our following a decision procedure that rules out harming an innocent person will prevent us from doing the act that would produce the best consequences. Similarly, there will be some circumstances in which stealing, breaking our promises, etc., would produce the best consequences. Still, our following a decision procedure that generally rules out such acts will – in the long run and on the whole – produce better consequences than our making act-utilitarian calculations on a case-by-case basis.

But now act-utilitarianism is in a seemingly paradoxical position. On the one hand, the theory holds that an act is morally permissible if and only if there is no alternative act whose consequences have greater expected utility. In its account of moral permissibility, act-utilitarianism makes no reference to rules about killing, promise breaking, attending to welfare of others to whom one has special connections, etc. On the other hand, act-utilitarianism tells agents to make their day-to-day moral decisions by following such rules. To illustrate the paradox, we should imagine an agent who follows such rules by,

for example, keeping a promise or refusing to steal. This agent might well feel confused when told that his act was nevertheless impermissible because, in the particular circumstances he faced, breaking the promise or stealing *would* have produced a bit more utility. Our agent might reasonably ask, 'Am I supposed to follow the optimific decision procedure and then do the act thus selected, or am I supposed to abandon the decision procedure you told me to follow and instead do an act selected by a different decision procedure, one which you told me *not* to follow?'

There is a mushrooming literature on what an act-utilitarian can say about the assessment of an agent who follows the prescribed decision procedure but does not select the expected-utility-maximizing act.⁴ I will not investigate that matter more here. I merely note that act-utilitarianism has some explaining to do, since the instructions it gives to the agent seem to be contradictory.

Where act-utilitarianism and rule-utilitarianism most closely agree is in their approach to which laws and to social policies there should be. I think most act-utilitarians and rule-utilitarians would agree that the law should allow euthanasia, that there should be a default presumption that people agree to have their organs donated once their brain dies, that tax policies should incentivize recycling, travel by land rather than air, and investment in renewable energy. On many social issues, various utilitarians might disagree with one another about what conclusion to draw, because they disagree on the probabilities of different benefits and harms. But their way of assessing social issues is the same. At the level of assessing actual or possible policies, laws, and rules, there is no significant difference between act-utilitarianism and rule-utilitarianism.

As I indicated, act-utilitarianism and rule-utilitarianism disagree about what makes acts morally required or wrong. For act-utilitarianism, what makes acts morally required is a fact about how their consequences compare with the consequences of alternative acts. For rule-utilitarianism, what makes acts morally required is a fact about how the consequences of rules that require them compare with the consequences of alternative rules.

⁴ The literature started mushrooming in response to Parfit's *Reasons and Persons*, 31–40. See (for example) Bart Streumer, 'Can Consequentialism Cover Everything?', *Utilitas* 15 (2003): 247–57 and Gerald Lang, 'A Dilemma for Objective Act-consequentialism', *Politics, Philosophy, and Economics* 3 (2004): 221–39. A recent addition to this literature by one of the most sophisticated and eloquent writers on the subject is Philip Pettit, 'The Inescapability of Consequentialism', in Ulrika Heuer and Gerald Lang (eds), *Luck, Value, and Commitment: Themes from the Ethics of Bernard Williams* (Oxford: Oxford University Press, 2012), 41–70.

As against rule-utilitarianism, act-utilitarians point out that, in order for rule-utilitarianism to avoid extensional equivalence with act-utilitarianism, rule-utilitarianism must hold at least some possible acts to be either wrong or merely optional that act-utilitarianism holds to be morally required. And, if rule-utilitarianism holds at least some possible acts to be wrong or merely optional that act-utilitarianism holds to be morally required, rule-utilitarianism cannot accept that all expected-utility-maximizing acts are morally required. Of course, to act-utilitarians, rule-utilitarianism is mistaken to accept that the set of *expected-utility-maximizing* acts and the set of *morally required* acts might diverge.

However, in denying that all expected-utility-maximizing acts are morally required, rule-utilitarianism is on the side of common-sense morality. For, in some cases, maximizing expected utility seems either morally wrong or merely optional. What such cases illustrate is that we intuitively think that there are various moral constraints, moral duties of loyalty, and moral prerogatives that can conflict with act-utilitarianism's requirement that one always maximize expected utility.

The moral constraints are prohibitions on killing or physically harming the innocent, stealing or destroying others' property, breaking promises, telling lies, and threatening to do any of these things. These prohibitions apply even where doing one of these acts would produce somewhat greater utility. Intuitively, it certainly seems that killing an innocent person when this would increase expected aggregate welfare a little would be morally wrong. Likewise, for stealing or destroying others' property, breaking promises, etc.

In cases where doing one of these acts would prevent *enormous* aggregate harm, we might have much less confidence that on balance the best thing to do would be to comply with the prohibition. But let's steer away from the less certain cases and back to the more certain ones. And the more certain cases are ones in which violating the prohibition would produce only *a little* more expected aggregate utility. These cases definitely are a problem for act-utilitarianism.

The moral duties of loyalty are special obligations you have to those with whom you have certain special connections, such as you do to your lover, friends, and family (at least normally). To be more specific, when deciding how to allocate of your own time, energy, and other resources, you should give some (but not infinite) priority to the welfare of your family and friends. There are also special obligations of gratitude to your benefactors and obligations of apology and compensation to those whom you wrongly harmed.

These duties of loyalty and special obligations of gratitude, apology, and compensation can conflict in individual cases with the requirement to maximize expected aggregate utility impartially calculated. Imagine that you could give your kidney to your forty-year-old sister or instead to a twenty-year-old stranger. There are many reasons you might think that expected aggregate utility would in fact be maximized if your kidney went to the twenty-year-old stranger rather than to your sister. Sometimes the difference in what is at stake for various parties will suggest that you should indeed override your duty of loyalty to your sister (suppose that saving the twenty-year-old stranger would somehow prevent a world war). Nevertheless, you have a strong duty of loyalty to your sister that, unless the case is extreme, outweighs your duty to do good for others in general.

Another element of common-sense morality is the idea that, even where no duties of loyalty, gratitude, apology, or compensation to others come into play, you are not morally required to be strictly impartial in deciding what to do with your own time, energy, and other resources. Admittedly, sometimes morality requires you to decide what to do with complete impartiality as between your own good and the equal good of any other person. For example, you might be the fire chief and have to decide to which side of the burning town to dispatch the fire trucks. And, sometimes, total disregard of your own good is mandatory. For example, when you are deciding what grade to assign to a student's work, you should be completely insensitive to possible benefits or harms to you. Nevertheless, sometimes you are deciding what to do with your own resources and you have the prerogative not to be impartial. For example, the world would be a bit better place if you devote your Sunday afternoons to visiting lonely people in the local rest home than if you spend your Sunday afternoons refining your amateur-level artistic or athletic skills. But, as Samuel Scheffler's *The Rejection of Morality* stressed, you have a moral prerogative to choose benefits for yourself even when you could instead have chosen somewhat larger benefits for others.⁵

The qualifier 'somewhat' is definitely necessary. The prerogative to be partial towards yourself has limits. In other words, the gap between the size of the benefit to yourself and the size of the aggregate benefit to others can be big enough to require you to sacrifice your own good for the sake of benefiting others. It isn't morally permissible to choose a trivial benefit for yourself over an equally probable massive benefit for others.

⁵ Samuel Scheffler, *The Rejection of Consequentialism* (Oxford: Oxford University Press, 1982), chs. 2–3.

Where is the line between permissible use of the prerogative and excessive self-interest? Perhaps the answer is indeterminate. Definitely there is no consensus on where the line is. Despite these important difficulties, I do think Scheffler was correct that there is a widely shared sense both that there is some degree of prerogative and that there is some limit to it.

Equally important is Michael Slote's observation that common-sense morality takes the prerogative to allow you not only to give your own good somewhat *greater* weight than the good of others when you are deciding what to do with your own time, energy, or other resources, but also to give your own good somewhat *less* weight.⁶ Imagine that you deliberately impose a loss on yourself for the sake of preventing a smaller loss to someone else. In deliberately not maximizing aggregate welfare here, you haven't done anything morally wrong. Indeed, most of us think that this other person should be very grateful and that everyone who hears of your act should commend your altruism.

In addition to the moral constraints, duties of loyalty, and moral prerogatives just outlined, common-sense morality also contains a general moral requirement to do good for others. Act-utilitarianism gets into conflict with the moral constraints, duties of loyalty, and moral prerogatives far more often than it does with the general moral requirement to do good for others. Indeed, the only time when the general duty to do good for others clearly conflicts with act-utilitarianism is when doing good for others would impose a loss on yourself that is larger than the combined benefits for others.

Singer's work mentions most of the moral duties and prerogatives I've listed. That is wise. For if he completely ignored these constraints, duties of loyalty, and moral prerogatives, his discussions would immediately provoke such questions as 'what about the moral constraint on killing the innocent?' or 'what about the moral duty to attend first to the affects on one's own children?' or 'what about the moral prerogative to give somewhat more weight to one's own welfare?'

However, in very many moral contexts, the constraints just do not come into play, or duties of loyalty do not come into play, or the prerogative to give one's own good more or less weight than the equal good of others does not come into play. For example, the constraints mentioned above are not much relevant to questions of population policy. Duties of loyalty are not prominent in questions about the environment. And the prerogative to give

⁶ Michael Slote, *Common-sense Morality and Consequentialism* (London: Routledge and Kegan Paul, 1985), ch. 1, and *From Morality to Virtue* (Oxford: Oxford University Press, 1992), 13–16.

one's own good greater or lesser weight in one's thinking does not come into play in any context in which one's own good would not be affected by the decision.

On the other hand, many of the contexts in which Singer has made significant contributions are ones where moral constraints might well be relevant. The constraint on killing the innocent is potentially relevant to killing animals, human fetuses, and those who request death. Moreover, the prerogative to give one's own good greater weight than the equal good of others will very often come into play, and will be especially salient in the two areas where the demandingness of morality is most prominent – see chapters 8 and 9 of Singer's *Practical Ethics*. Still, as I have said, there are other contexts in which no constraints, duties of loyalty, or prerogatives come into play, and yet we feel that the context is one in which moral issues are definitely at stake.

As indicated earlier, act-utilitarians typically think that there is an impartial, utilitarian rationale for people's internalizing some moral constraints, duties of loyalty, and moral prerogatives (though act-utilitarians are typically reformers and thus advocate internalization of *improved* moral constraints, duties of loyalty, and moral prerogatives). If people have internalized some moral constraints, if they feel especially responsible for the welfare of their family and friends, and if they feel free in many contexts to focus on their own welfare or alternatively to ignore their own welfare, more overall welfare will result than if people feel they must calculate expected aggregate utility all the time and then rigorously comply with the conclusions mandated by their calculations. And some especially engaging parts of Singer's work directly address the question of which constraints, duties of loyalty, and prerogatives are the ones whose internalization would maximize expected utility. He also directly addresses the question of which rules are the ones whose public endorsement would maximize expected utility.

The question of which constraints, duties of loyalty, and prerogatives are the ones whose internalization would maximize expected utility is irresistible when we turn to practical policies, and especially irresistible when we turn to some new moral problem created by the development of some new technology. Notice the rule-utilitarian aspects of this question. First, the evaluation is of different possible codes of moral rules (which typically will contain constraints, duties of loyalty, and prerogatives). Second, these different possible codes are evaluated in terms of the consequences of the codes' being taught to and then collectively accepted by whole generations of moral agents (not merely in terms of internalization by some individual agent or only by some elite group). Third, this evaluation is from a

manifestly impartial point of view, a point of view that accords equal importance to the welfare of each individual.

Because this approach conceives of morality as a collective endeavour with an attractive impartial purpose, i.e. promotion of aggregate welfare, this approach seems to me intuitively attractive *in itself*, that is, intuitively attractive even before we consider the plausibility of the rules it ends up endorsing. I also think, albeit more tentatively, that this approach ends up endorsing intuitively plausible rules. The *overall* intuitive plausibility of the approach is made up of the intuitive plausibility of its starting assumptions and the intuitive plausibility of its implications. In other words, the philosophical argument for rule-utilitarianism seems to me best understood as what John Rawls called a reflective equilibrium argument.⁷

I did not employ a utilitarian premise in my argument for rule-utilitarianism. To employ a utilitarian premise in an argument for rule-utilitarianism is to beg the question whether utilitarianism is a good starting place. And I am not arguing for rule-utilitarianism via religious premises. I cannot do that, since I am not religious.

Nevertheless, I do think rule-utilitarianism is the natural position for many religious believers and especially for Christians to hold. A familiar doctrine is that God chooses the rules whose internalization by human beings would produce the greatest good and human beings' job is to follow these rules. Bishop Berkeley wrote, 'In framing the general laws of nature, it is granted we must be entirely guided by the public good of mankind, but not in the ordinary moral actions of our lives . . . The rule is framed with respect to the good of mankind; but our practice must be always shaped immediately by the rule.'⁸

In some ways, rule-utilitarianism seems easier to defend on theist rather than atheist grounds. One of the problems rule-utilitarianism faces is the problem of predicting the consequences of internalization of various possible rules.⁹ Of course there is no difficulty in predicting that the consequences of any code of rules that allows wife-beating would be terrible. But would internalization of a rule that allows human genetic engineering have better consequences than internalization of a rule that forbids human genetic engineering? Reasonable people can disagree about this. Indeed, a

⁷ John Rawls, 'Outline for a Decision Procedure in Ethics', *Philosophical Review* 60 (1951): 177–97; *A Theory of Justice* (Cambridge, MA: Harvard University Press, 1971), 19–21, 46–51.

⁸ George Berkeley, *Passive Obedience, or, the Christian Doctrine of Not Resisting the Supreme Power, Proved and Vindicated upon the Principles of the Law of Nature* (1712), section 31.

⁹ I discuss this matter at length in my *Ideal Code, Real World: A Rule-consequentialist Theory of Morality*, (Oxford: Oxford University Press, 2000), ch. 5.

general objection to rule-utilitarianism is that often people cannot know which rules are the ones whose internalization would have the very best consequences, and don't even have the information about possible consequences and their probabilities needed to calculate expected utility. However, if an *all-knowing* and perfectly good God chooses certain rules for us to follow, then presumably the lack of knowledge problem is solved.

But does God really offer us a way out of rule-utilitarianism's epistemological struggles? Will thinking about the rules reported to have come from God help us ascertain which rules are the ones rule-utilitarianism endorses? Difficulties in interpreting the rules reported to have come from God have been apparent since the rules were first reported. Furthermore, even if we could be certain exactly what God meant for humans to do 2,000 years ago, why think the best rules for today would be the same as the best rules for 2,000 years ago?

Most religious believers make every effort to attribute to God commands that correspond to rules that are morally sensible for human beings. What distinguishes rules that are morally sensible for human beings from rules that are not morally sensible for human beings? Well, one criterion that jumps to mind is the rule-utilitarian one: morally sensible rules for human beings are the ones whose internalization by human beings would produce the best consequences, impartially considered.

There is no denying, however, that some of the commands reported long ago to have come from God seem very unlikely to mirror rules that meet the rule-utilitarian criterion. This fact seems to me to leave us with the following possible alternative responses:

Alternative 1. These reports from long ago about God's commands were mistaken.

Alternative 2. Even if these reports from long ago weren't mistaken, the commands in question were *not* meant to apply outside of a long-gone context.

Alternative 3. Even if these reports from long ago weren't mistaken, and even if all God's commands *were* meant to apply in all (even modern) contexts, and even if some of God's commands *appear* to enjoin rules that do not meet the rule-utilitarian criterion, this appearance is an illusion: all God's commands do in fact enjoin rules that meet the rule-utilitarian criterion.

Alternative 4. These reports from long ago weren't mistaken. Furthermore, all God's commands were meant to apply in all contexts. Moreover, some of God's commands *really do* enjoin rules that do not meet the rule-utilitarian criterion. So God is not a rule-utilitarian.

Which of these alternatives is most plausible? Of course, atheists insist on alternative 1. They deny the existence of God and thus deny the existence of God's commands. Many religious believers will also accept alternative 1. Indeed, I cannot see how someone who believes in a perfectly good God could believe that God ever made some of the commands reported long ago.

Turn now to religious believers who think that authoritative religious texts supply credible evidence that God did indeed make those commands. These religious believers reject alternative 1. So which of alternatives 2, 3, and 4 do these religious believers accept? The least plausible of these alternatives is alternative 3. Some of the commands reported to have been made by God cannot plausibly be thought really to correspond to rules that meet the rule-utilitarian criterion. Hence I think those religious believers who do not disbelieve the reports from long ago about God's commands face a choice between alternatives 2 and 4.

Therefore, which is the more plausible of the following two alternatives: that even if the reports from long ago about God's commands weren't mistaken, the commands that seem incompatible with a rule-utilitarian criterion of correct moral rules were not meant to apply outside of a long-gone context, or that God was not a rule-utilitarian? When we think about how silly or severe some of these reported commands were, the most plausible explanation is that these reports from long ago about God's commands were mistaken. The second most plausible explanation is that the commands were not meant to apply outside of a long-gone context. Much less plausible than those two explanations are the two explanations that the commands in question are somehow compatible with rule-utilitarianism and that God was not a rule-utilitarian.

I have been discussing the idea that, if there is an all-knowing God, rule-utilitarianism is easier to defend against the objection that we cannot know which rules are the ones whose internalization would produce the best consequences than it is if there is no such God. Let me now turn to the idea that rule-utilitarianism is easier to defend than act-utilitarianism is if an atheistic world-view is accepted than if a traditional Christian world-view is accepted.

Act-utilitarianism is regularly attacked for being excessively demanding. The normal form of this objection is that the theory requires a level of self-sacrifice from individuals for the sake of aiding others that is unreasonable to require. If we accept an atheistic world-view, individuals are compensated for whatever sacrifices they make for the sake of others *only* either by the pleasure or sense of achievement that the individual gets from these acts, or by admiration and gratitude from other people. The rewards that come

from within the individual or from other people are unlikely to compensate for the sacrifices that act-utilitarianism requires of the individual. Indeed, most atheists think act-utilitarianism requires so much sacrifice that its demands are seriously unreasonable.

If there is an all-powerful and perfectly good God, on the other hand, then perhaps the sacrifices that act-utilitarianism calls for are fully compensated by God. The impact of divine rewards for the kind and divine punishments for the unkind extinguishes the demandingness objection to act-utilitarianism.¹⁰ In this respect, act-utilitarianism is far easier to defend if there is an all-powerful and perfectly good God than if there is not.

To those atheists who find act-utilitarianism excessively demanding, rule-utilitarianism can look more attractive than act-utilitarianism. Rule-utilitarianism's strictures about benefiting others are normally less demanding than act-utilitarianism's. Rule-utilitarianism starts with the question 'which rule about aid is the one whose internalization by us as a group has the greatest expected utility?', rather than the question 'which act can I as an individual do now to maximize expected utility?'

Whether the duty to aid that rule-utilitarianism specifies is, though not excessively demanding, intuitively plausible is a further matter.¹¹ Most problematic for rule-utilitarianism are cases where one knows that, because others are not doing what is required of them to help, one could produce much more aggregate utility by doing more, indeed far more, than one's own fair share. The classic examples of this are the cases where others in a position to help are not contributing to some rescue effort and one could save more lives by contributing more than one's fair share. In such cases, Singer quickly moves in an act-utilitarian direction, despite the extreme demands on him that this entails. Here he sets an example no one could fail to admire and only a few manage to match.

In conclusion, the philosophical threads out of with Singer's morality is woven are a rule-utilitarian approach to questions about which laws and policies and rules we should have and an act-utilitarian willingness to undertake additional sacrifices to make up for others' failing to do their fair share. On reflection, we might think that this act-utilitarian element renders the morality fit for saints rather than for ordinary human beings. We can hardly deny, however, its appeal as an ideal.

¹⁰ Sidgwick, *Methods of Ethics*, 503–6.

¹¹ See my *Ideal Code, Real World*, ch. 8 and 'Rule-consequentialism and Internal Consistency: A Reply to Card', *Utilitas* 19 (2007): 514–19.