

Penultimate draft: for published version see *Analysis*, October 1990, pp. 253-56.

## CONSEQUENTIALISM AND THE UNFORESEEABLE FUTURE

By Alastair Norcross

If consequentialism is understood as claiming, at least, that the moral character of an action depends only on the consequences of the action, it might be thought that the difficulty of knowing what all the consequences of any action will be poses a problem for consequentialism. J. J. C. Smart writes that in most cases

...we do not...need to consider very remote consequences, as these in the end approximate rapidly to zero like the furthest ripples on a pond after a stone has been dropped into it. ... The necessity for the 'ripples in the pond' postulate comes from the fact that usually we do not know whether remote consequences will be good or bad. Therefore we cannot know what to do unless we can assume that remote consequences can be left out of account. ... If [such a postulate] is not accepted, not only utilitarianism, but also deontological systems like that of Sir David Ross, who at least admits beneficence as one prima facie duty among others, will be fatally affected. (J. J. C. Smart, An Outline of a System of Utilitarian Ethics, in Smart and Williams, Utilitarianism for and against, Cambridge: Cambridge University Press, 1973, 33-34.)

Smart's claim seems to be that we needn't worry too much about the unforeseeable future, because most of our choices make very little difference to the state of the universe beyond what we can reasonably foresee. I do not see any good reason to believe Smart's

'ripples in the pond' postulate, but I also do not think that consequentialism needs such a postulate.

Smart's worry seems to be that our ignorance of remote consequences will strip us of the ability to make rational choices on consequentialist grounds. After all, if my choice now will go on making a difference to the world far beyond what I can foresee, how can I be justified in my belief that I am doing the right thing? I think that the consequentialist can admit that we are often faced with choices concerning which we have very little reason to believe that one choice will have better total consequences than the other. I also think that it is possible to make rational decisions in such cases on consequentialist grounds alone. Consider the following illustration: you are visiting a totalitarian dictatorship on a humanitarian mission, when the dictator himself, whom you have reason to trust, informs you that the fates of two political prisoners of conscience, Smith and Jones, are in your hands. You can specify whether Smith lives or dies. If you do not specify either fate for Smith, you will be taken to have specified death. In addition, you are told that the fate of Jones is tied to your decision about Smith. At this very moment, in a different part of the presidential palace, a coin is being tossed to link your decision about Smith to a decision about Jones. If the coin lands heads, life for Smith will also mean life for Jones, and death for Smith will mean death for Jones. If the coin lands tails, life for Smith will mean death for Jones, and death for Smith will mean life for Jones. You will not learn of the result of the coin toss, however. If your only concern is with the fates of Smith and Jones, it is clear what your choice must be. If you specify life for Smith, it will be certain that he will live, and it will be as likely as not that Jones will live too. If you specify death for Smith (even by a refusal to take part in the

dictator's scheme), it will be certain that he will die, and it will be as likely as not that Jones will die too. You can, of course, employ rational choice theory, discounting the various possible outcomes by their probabilities. On this basis, a choice of life for Smith will have an expected outcome of one and a half lives saved; whereas a choice of death for Smith will have an expected outcome of half a life saved. If your only concern is with the fates of Smith and Jones, you must specify life for Smith.

What happens if we add a third prisoner, Wilson, whose fate is also tied to your decision about Smith by another coin toss? If you specify life for Smith, it will be certain that Smith will live, as likely as not that Jones will live, and as likely as not that Wilson will live (not as likely as not that both will live). Rational choice theory gives the result that a choice of life for Smith has an expected outcome of two lives saved, whereas a choice of death for Smith has an expected outcome of one life saved. Again, it is clear that you must choose life for Smith. Notice, however, that the three-person case has an important feature which the two-person case lacks. In the two-person case, a choice of life for Smith could not have a worse outcome than a choice of death for Smith. In the three-person case, it is possible, though unlikely, that a choice of life for Smith will result in fewer lives saved than a choice of death for Smith. This should not affect the rationality of your choice. If you have no more reason to think that a choice of death for Smith will benefit Jones and Wilson than a choice of life for Smith, and you know which will benefit Smith, you choose life for Smith. In this case, the unforeseeable consequences of your choice are greater than the foreseeable consequences, but the basis of your choice is still clear.

I do not think that your choice is any less clear when there are seven million other prisoners whose lives depend on your decision for Smith. Why would anyone think otherwise? There may be a temptation to regard one life as trivial when compared with seven million. What difference will a choice of life or death for Smith make when compared with the millions who will surely die whatever you choose? Or perhaps we could say that it is not so much that one more life is trivial compared with several million, but rather that morality should not have anything to say about such a difference. Bernard Williams could be taken to be describing such a view when he talks of a moral agent for whom

...there are certain situations so monstrous that the idea that the processes of moral rationality could yield an answer in them is insane: they are situations which so transcend in enormity the human business of moral deliberation that from a moral point of view it cannot matter any more what happens. (Bernard Williams, A Critique of Utilitarianism, in Smart and Williams, p. 92.)

Williams contrasts such a view with consequentialism, which 'will have something to say even on the difference between massacring seven million, and massacring seven million and one.' (p. 93) One can certainly sympathize with the agent who is so horrified at the scale of a massacre that she finds it difficult to deliberate rationally in the circumstances. This does not, however, support the view that from a moral point of view it cannot matter anymore what happens. If there really is no moral difference between massacring seven million and massacring seven million and one, the allied soldier arriving at Auschwitz can have no moral reason for preventing the murder of one last Jew before the Nazi surrender. The Nazi himself can have no moral reason for refraining

from one last murder. While Williams's moral agent is berating the universe for transcending the bounds of rationality, the consequentialist is saving a life. It is not hard to guess which of these agents I would rather have on my side.

The example of the difference between massacring seven million and massacring seven million and one is importantly different from my example of a choice of life or death for Smith tied to the fates of seven million other prisoners. Williams's example does not pose a problem of application for the consequentialist. It is clear that seven million and one deaths are worse than seven million. Rather, Williams's example is supposed to illustrate a feature of consequentialist reasoning which Williams finds disturbing. My example might be thought to be problematic for the consequentialist, because the agent has very little reason to think that a choice of life for Smith will have overall better consequences than a choice of death for Smith. I think, however, that there is an important similarity between the two cases. In each case there may be a temptation to abandon consequentialist methods of reasoning on the grounds that what one is able to bring about, or able to be certain to bring about, is somehow dwarfed by the scale of what else is beyond one's control, or beyond one's epistemic reach. In each case, the cure is to focus on what is within one's control. Whether one person lives or dies is always morally significant, indeed it is this very significance which is at the root of our horror at the sort of massacres described by Williams. If one can decide whether Smith lives or dies, and there is no more reason to believe that life for Smith will have consequences for others which are worse than the consequences of death for Smith, the choice is just as clear in the seven million and one person case as in the two person case.

My example, of course, does not exactly model our epistemic situation with regard to the future. We do not operate only with subjective probabilities of one and of one half. A more complicated example could be devised as a variation on the situation I describe. My point would remain. It is rational for a consequentialist to accept that a choice with foreseeable good consequences may, for all she knows, have unforeseeable bad consequences which massively outweigh the good consequences, and yet to make the choice purely on the basis of the foreseeable consequences.<sup>1</sup>

Syracuse University,

Syracuse, NY 13244-1170, U.S.A.

---

<sup>1</sup> This paper has benefitted from discussions with Jonathan Bennett, Joel Kidder, Steven Lee and Rick Wiley.