



Decision-Theoretic Consequentialism and the Nearest and Dearest Objection

Author(s): Frank Jackson

Source: *Ethics*, Vol. 101, No. 3 (Apr., 1991), pp. 461-482

Published by: [The University of Chicago Press](#)

Stable URL: <http://www.jstor.org/stable/2381464>

Accessed: 15-09-2015 15:44 UTC

---

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact [support@jstor.org](mailto:support@jstor.org).



The University of Chicago Press is collaborating with JSTOR to digitize, preserve and extend access to *Ethics*.

<http://www.jstor.org>

# Decision-theoretic Consequentialism and the Nearest and Dearest Objection\*

*Frank Jackson*

Our lives are given shape, meaning and value by what we hold dear, by those persons and life projects to which we are especially committed. This implies that when we act we must give a special place to those persons (typically our family and friends) and those projects. But, according to consequentialism classically conceived, the rightness and wrongness of an action is determined by the action's consequences considered impartially, without reference to the agent whose actions they are consequences of. It is the nature of any particular consequence that matters, not the identity of the agent responsible for the consequence. It seems then that consequentialism is in conflict with what makes life worth living. I take this to be one part of Bernard Williams's well-known attack on consequentialism.<sup>1</sup>

One way to reply to it would be to break the implicit connection between acting morally and living a life worth living. Doing what is morally right or morally required is one thing; doing what makes life worth living is another. Hence, runs the reply, it is no refutation of a moral theory that doing as it enjoins would rob life of its shape and meaning.

This is a chilling reply and I will say no more about it. My reply will be that consequentialism—properly understood—is perfectly compatible with the right actions for a person being in many cases actions directed toward achieving good consequences for those persons and projects that the agent holds dear. Consequentialism, I will argue, can make plausible sense of the moral agent having and giving expression in action to a special place for family, friends, colleagues, chosen projects, and so on and so forth.

\* I am indebted to discussions with a number of audiences and to comments from Michael Smith, Peter Singer, Philip Pettit, and a referee (to whom I owe the title).

1. In, e.g., Bernard Williams, "A Critique of Utilitarianism," in J. J. C. Smart and Bernard Williams, *Utilitarianism: For and Against* (Cambridge: Cambridge University Press, 1973).

*Ethics* 101 (April 1991): 461–482

© 1991 by The University of Chicago. All rights reserved. 0014-1704/91/0103-0482\$01.00

I will start by explaining how, in my view, consequentialism should be understood. This explanation will help us sort out certain potential confusions in addition to providing the springboard for our reply to the nearest and dearest objection.

I should emphasize that I claim no great originality for my account of consequentialism. It is, I think, a natural extension of what, for instance, J. J. C. Smart had in mind all along (shorn of the commitment to a utilitarian construal of consequences), though he did not say it quite the way I will.<sup>2</sup> I hope that my way of putting things will make certain matters clearer.

## UNDERSTANDING CONSEQUENTIALISM

Consequentialism approaches the question of whether an action is right or wrong in terms of a comparison of the possible outcomes of the action with the possible outcomes of each available alternative to that action. The notion of a possible outcome of an action is interpreted so as to include the action itself, and the comparison of the various outcomes is carried out in terms of a consequentialist value function. The interesting question of exactly what makes a value function warrant being described as consequentialist can here be left to one side. The details of the value function will not particularly concern us; any reasonable ranking of outcomes of the usual agent neutral kind will serve our purposes in what follows. Similarly, exactly how the available alternatives to the action in question are specified can be left vague. What will, however, concern us is how the values assigned to the outcomes feed into the determination of what ought to be done. We will be presupposing that the matter is approached in the usual maximizing way—classical consequentialism is our subject, not satisficing varieties thereof<sup>3</sup>—but that in itself leaves a major issue open, an issue which will turn out to be crucial for the argument of the article. This major issue can be most easily approached via a simple example.

### *The Drug Example, Mark 1*

Jill is a physician who has to decide on the correct treatment for her patient, John, who has a minor but not trivial skin complaint. She has three drugs to choose from: drug A, drug B, and drug C. Careful con-

2. J. J. C. Smart, "An Outline of a System of Utilitarian Ethics," in Smart and Williams. It is arguable that something like the account can also be found in some classical presentations, e.g., Jeremy Bentham, *An Introduction to the Principles of Morals and Legislation* (London: Athlone, 1970); and Henry Sidgwick, *The Methods of Ethics*, 7th ed. (Chicago: University of Chicago Press, 1907), though other interpretations are very possible, and, as we will see, a much quoted passage from Sidgwick points in a quite different direction.

3. A satisficing variety is expounded in Michael Slote, *Commonsense Morality and Consequentialism* (London: Routledge & Kegan Paul, 1985), as worthy of serious attention. The discussion is set in the context of a general defense of satisficing, as opposed to optimizing or maximizing, approaches to decision theory.

sideration of the literature has led her to the following opinions. Drug A is very likely to relieve the condition but will not completely cure it. One of drugs B and C will completely cure the skin condition; the other though will kill the patient, and there is no way that she can tell which of the two is the perfect cure and which the killer drug. What should Jill do?

The possible outcomes we need to consider are: a complete cure for John, a partial cure, and death. It is clear how to rank them: a complete cure is best, followed by a partial cure, and worst is John's death. That is how Jill does, and also how she *ought* to, rank them. But how do we move from that ranking to a resolution concerning what Jill ought to do? The obvious answer is to take a leaf out of decision theory's book and take the results of multiplying the value of each possible outcome of each contemplated action by Jill's subjective probability of that outcome given that the action is performed, summing these for each action, and then designating the action with the greatest sum as what ought to be done. In our example there will be three sums to consider, namely:

$$\begin{aligned} Pr(\text{partial cure/drug A taken}) \times V(\text{partial cure}) \\ + Pr(\text{no change/drug A taken}) \\ \times V(\text{no change}); \end{aligned}$$

$$\begin{aligned} Pr(\text{complete cure/drug B taken}) \times V(\text{complete cure}) \\ + Pr(\text{death/drug B taken}) \\ \times V(\text{death}); \text{ and} \end{aligned}$$

$$\begin{aligned} Pr(\text{complete cure/drug C taken}) \times V(\text{complete cure}) \\ + Pr(\text{death/drug C taken}) \\ \times V(\text{death}). \end{aligned}$$

Obviously, in the situation as described, the first will take the highest value, and so we get the answer that Jill should prescribe drug A. The obvious answer all along. The difference between a complete cure and a partial cure in the case of a minor skin complaint does not compensate for a significant risk of death, as we might say it in English.

Generalizing, the proposal is to recover what an agent ought to do at a time according to consequentialism from consequentialism's value function—an assignment of value that goes by total consequent happiness, average consequent preference satisfaction, or whatever it may be in some particular version of consequentialism—together with the agent's subjective probability function at the time in question in the way familiar in decision theory, with the difference that the agent's preference function that figures in decision theory is replaced by the value function of consequentialism. That is to say, the rule of action is to maximize  $\sum_i Pr(O_i/A_j) \times V(O_i)$ , where  $Pr$  is the agent's probability function at the time,  $V$  is consequentialism's value function,  $O_i$  are the possible outcomes, and

$A_j$  are the possible actions. We can express the idea in English by saying that whereas decision theory enjoins the maximization of expected utility, consequentialism enjoins the maximization of expected *moral* utility.<sup>4</sup>

How else might one seek to recover what, according to consequentialism, a person ought to do from consequentialism's value function? Two alternatives to the decision-theoretic approach call for discussion.

We can think of consequentialism's value function as telling us what, according to consequentialism, we ought to desire. For a person's desires can be represented—with, of course, a fair degree of idealization—by a preference function which ranks state of affairs in terms of how much the person would like the state of affairs to obtain, and we can think of consequentialism as saying that the desires a person ought to have are those which would be represented by a preference function which coincided with consequentialism's value function. The other ingredient in the decision-theoretic account of what consequentialism says a person ought to do, the agent's subjective probability function, is an idealization of the agent's beliefs. Hence, the decision-theoretic account is one in terms of what the person ought to desire and in fact believes. But in addition to distinguishing what a person in fact desires from what he or she ought to desire, we also distinguish what a person in fact believes from what a person ought to believe. And in a sense of 'ought' which has a moral dimension—there is, for instance, such a thing as culpable ignorance. Hence, it might well be suggested that we should recover the consequentialist answer to what a person ought to do from the value function via what that person ought to believe rather than from what he or she in fact believes.<sup>5</sup>

However, the clearest cases of culpable ignorance can be handled in terms of what a person in fact believes. The decision problem which faces a doctor considering whether to prescribe a certain drug is not simply the choice between prescribing the drug and not doing so—though we may pretend that it is that simple in order to make some point that is independent of the complexities—it is more accurately described as the choice between: deciding now to prescribe the drug, deciding now against prescribing the drug, and postponing the decision until more information has been obtained and, on the basis of what one then knows, deciding between prescribing the drug and not prescribing

4. Decision theory comes in a number of varieties. For example, in some  $Pr(O_i/A_j)$  is replaced by  $Pr(A_j \rightarrow O_i)$ . The points I wish to make here are independent of the particular variety. (Though I in fact favor the latter, which is indeed the most obvious way of capturing a *consequentialist* approach to matters provided the " $\rightarrow$ " is read appropriately.) For a recent discussion of the varieties, see Ellery Eells, *Rational Decision and Causality* (Cambridge: Cambridge University Press, 1982). Incidentally, the preference function in decision theory is often referred to as a value function, but I will reserve the latter term for what an agent ought to prefer in the moral sense.

5. Smart appears to favor a proposal of this kind. See also Philip Pettit and Geoffrey Brennan, "Restrictive Consequentialism," *Australasian Journal of Philosophy* 64 (1986): 438–55.

the drug. Now, in the same way that prescribing the drug and not prescribing the drug have an expected moral utility, so does obtaining more information and doing what then has greatest expected moral utility; and we can investigate the conditions under which getting more information and doing what then has greatest expected moral utility has itself greater expected moral utility than either prescribing the drug or not prescribing the drug. It is easy to prove the following. Getting more information and then doing what has greatest moral utility has itself greatest moral utility provided the possible change in utility consequent on the new information when weighed by the probability of getting that new information is great enough to compensate for the effort and cost of getting the new information.<sup>6</sup> Thus, working solely with a person's subjective probability function, with what he or she actually believes, we can distinguish plausibly between cases where more information ought to be obtained and where we may legitimately rest content with what we have. Hence, it seems to me at least arguable that our approach to what a person ought to do according to consequentialism in terms of what he or she ought to desire and does in fact believe does not need to have the reference to what is believed replaced by a reference to what ought to be believed.<sup>7</sup> However, the bulk of what I have to say about the nearest and dearest objection to consequentialism is independent of this issue.

The other possible account of how to recover what a person ought to do from consequentialism's value function that we need to consider holds that a person's beliefs, rational or not, do not come into the picture. What is crucial is simply which action *in fact* has, or would have, the best consequences.<sup>8</sup> Many consequentialists write as if this was their view. In a well-known passage Sidgwick says "that Universal Happiness is the ultimate *standard* must not be taken to imply that Universal Benevolence is the only right . . . *motive* for action. . . . It is not necessary that the end which gives the criterion of rightness should always be the end at which we consciously aim."<sup>9</sup> Here it seems clear that he is assuming that what makes an act right—the criterion of rightness, as he puts it—is the extent to which it in fact achieves a certain end.<sup>10</sup> Similarly, Peter Railton dis-

6. For a clear presentation of the proof, see Paul Horwich, *Probability and Evidence* (Cambridge: Cambridge University Press, 1982), pp. 125–26.

7. For a more detailed development of this argument, see Frank Jackson, "A Probabilistic Approach to Moral Responsibility," in *Proceedings of the 7th International Congress of Logic, Methodology, and Philosophy of Science*, ed. R. Barcan Marcus et al. (Amsterdam: Elsevier, 1986), pp. 351–66.

8. Another approach in which the agent's beliefs do not come into the picture is one in which it is objective, one-place chances, rather than probabilities construed epistemically, of the various possible outcomes which matter, but I take the essentials of the critical discussion that follows to apply equally against this approach.

9. Sidgwick, p. 413.

10. And this is certainly how this passage is typically read; see, e.g., David O. Brink, *Moral Realism and the Foundations of Ethics* (Cambridge: Cambridge University Press, 1989), p. 257. And Brink is explicit in endorsing the idea that what makes an act right according

tinguishes *subjective consequentialism*, the doctrine that “whenever one faces a choice of actions, one should attempt to determine which act of those available would most promote the good, and should then try to act accordingly,” from *objective consequentialism*, “the view that the criterion of the rightness of an act or course of action is whether it *in fact* would most promote the good,” and goes on to argue in support of objective consequentialism.<sup>11</sup>

There are two problems with this proposal. First, it gives the intuitively wrong answer in the drugs case. In the drugs case, either it is prescribing drug B or it is prescribing drug C which is the course of action which would in fact have the best consequences—and Jill knows this, although she does not know which of the two it is—but neither prescribing drug B nor prescribing drug C is the right course of action for Jill. As we observed earlier, it is prescribing drug A which is the intuitively correct course of action for Jill despite the fact that she *knows* that it will *not* have the best consequences. We would be horrified if she prescribed drug B, and horrified if she prescribed drug C.

The second problem arises from the fact that we are dealing with an *ethical* theory when we deal with consequentialism, a theory about *action*, about what to *do*. In consequence we have to see consequentialism as containing as a constitutive part prescriptions for action. Now, the fact that an action has in fact the best consequences may be a matter which is obscure to an agent. (Similarly, it may be obscure to the agent what the objective chances are.) In the drugs example, Jill has some idea but not enough of an idea about which course of action would have the best results. In other examples the agents have very little idea which course of action would have the best results. This was the case until recently in the treatment of AIDS. Hence, the fact that a course of action would have the best results is not in itself a guide to action, for a guide

---

to consequentialism should be recovered from what in fact does or would happen. See also Fred Feldman, *Doing the Best We Can* (Dordrecht: Reidel, 1986).

11. Peter Railton, “Alienation, Consequentialism, and the Demands of Morality,” *Philosophy and Public Affairs*, vol. 13 (1984), reprinted in *Consequentialism and Its Critics*, ed. S. Scheffler (Oxford: Oxford University Press, 1988), pp. 93–133, p. 113, my emphasis. All citations to Railton will be to that in Scheffler, ed. Railton (p. 113, n. 24) mentions the decision-theoretic approach in passing, and it is unclear going on the printed word how much substantive disagreement there is between us on the question of what, according to consequentialism, a person ought to do. However, very helpful discussion with Railton (February 1990) has made it clear to me that we are in substantive disagreement. What is, in any case, clear from the printed word alone is that we are in substantive disagreement over how to answer the nearest and dearest objection, for in the footnote he remarks that his arguments go through independently of whether or not the decision-theoretic approach is adopted, whereas, as will become very clear, our treatment of the nearest and dearest objection rests quite crucially on the adoption of the decision-theoretic approach. In Bart Gruzalski, “The Defeat of Utilitarian Generalization,” *Ethics* 93 (1982): 22–38, a decision-theoretic approach is given a crucial role in assessing the relative merits of act and rule or generalized versions of utilitarianism.

to action must in some appropriate sense be present to the agent's mind. We need, if you like, a story from the inside of an agent to be part of any theory which is properly a theory in ethics, and having the best consequences is a story from the outside. It is fine for a theory in physics to tell us about its central notions in a way which leaves it obscure how to move from those notions to action, for that passage can be left to something which is not physics; but the passage to action is the very business of ethics.

Railton is well aware of the need to give an account of the passage to action, for he remarks that a "further objection [to objective consequentialism] is that the lack of any direct link between objective consequentialism and a particular mode of decision making leaves the view too vague to provide adequate *guidance* in practice." His reply is that "on the contrary, objective consequentialism sets a definite and distinctive criterion of right action, and it becomes an empirical question . . . which modes of decision making should be employed and when."<sup>12</sup> In short, Railton's proposal is, I take it, that the moral decision problem should be approached by setting oneself the goal of doing what is objectively right—the action that has in fact the best consequences—and then performing the action which the empirical evidence suggests is most likely to have this property.<sup>13</sup> However, this approach to the decision problem gives the wrong answers. I will illustrate the point with a modification of the drug example.

### *The Drug Example, Mark 2*

As before, Jill is the doctor and John is the patient with the skin problem. But this time Jill has only two drugs, drug X and drug Y, at her disposal which have any chance of effecting a cure. Drug X has a 90% chance of curing the patient but also has a 10% chance of killing him; drug Y has a 50% chance of curing the patient but has no bad side effects. Jill's choice is between prescribing X or prescribing Y. It is clear that she should prescribe Y, and yet that course of action is not the course of action most likely to have the best results, for it is not the course of action most likely to be objectively right. It has only a 50% chance of being objectively right, whereas prescribing drug X has a 90% chance of being objectively right.

This example is one among many. Consider, for instance, the question of whether or not to place a bet on a horse race. Clearly, it is often the right thing to do not to place a bet. There may be no horse about which the bookies are offering good enough odds. And yet in declining to place a bet you know that you are pursuing the one course of action guaranteed

12. Both passages are from Railton, p. 117, my emphasis.

13. See also his earlier remarks on objective consequentialism "not blurring the distinction between the *truth-conditions* of an ethical theory and its *acceptance-conditions*" (ibid., p. 116). Feldman explicitly takes this approach to the moral decision problem.



not to have the best outcome. Your problem, of course, is that although you know that there is a course of action with a better outcome, you do not know which one it is.

In general it seems to me potentially misleading to speak of consequentialism as giving the moral agent the aim of doing what has the best consequences. If it means that the agent ought to do what has the highest expected value or moral utility, where value is determined consequentially, then of course there is no problem. But it is easy to slide into thinking that consequentialism holds that people should aim at the best consequences in the sense of trying to select the option with the best consequences, whereas in fact most of the time we should select an option which we know for sure does not have the best consequences. Most of the time we are in the position of the person who declines to bet. The right option to select is a "play safe" one chosen in the knowledge that it does not have the best consequences and in ignorance of which option does have the best consequences.

I argued that consequentialism must, as I put it, tell a story from the inside about how to recover what an agent ought to do from consequentialism's value function, a story in terms of what is in the agent's mind at the time of action. I thus am agreeing with Thomas Nagel's claim that "morality requires of us not only certain forms of conduct but also the motives required to produce that conduct."<sup>14</sup> For the proposal I borrowed from decision theory meets this constraint, because it is in terms of the agent's probability function, that is, in terms of the agent's belief state at the time of action. Indeed, the proposal I borrowed from decision theory can be viewed as an account of what an agent ought to do that yields an account of what an agent's motives ought to be, of how an agent's mind ought to be as far as the springs of action go. For, as we remarked earlier, we can view consequentialism's value function as an account of how an agent's preference function, the agent's desires, ought to be—the preference function ought to assign the same values to the various states of affairs as does the value function. Hence, when the proposal recovers what an agent ought to do from the agent's probability function combined with consequentialism's value function, this can be described as recovering what an agent ought to do from what the agent believes combined with what the agent ought to desire and, thereby, as yielding a theory of right motivation.

Although consequentialism of the decision-theoretic kind described here (henceforth, consequentialism) has built into its very account of right action, a doctrine about right motivation, it is not committed to any particular view about the mental *processes* an agent ought to go through in deciding what to do; indeed it is compatible with consequentialism

14. Thomas Nagel, *The View from Nowhere* (Oxford: Oxford University Press, 1986), p. 191. See also Williams, "A Critique of Utilitarianism," p. 128, though the focus there is on utilitarianism rather than on ethics in general.

that at least sometimes one ought not to go through anything that might naturally be described as a *process* of thought at all. It is, for instance, compatible with consequentialism that an agent ought to go through a distinctively deontological cast of thought before acting. The agent may, for instance, have a terrible track record in calculating likely consequences and have found out by bitter experience that he does better following simple rules; or perhaps he knows that the world is under the control of a demon who rewards people for thinking in the Kantian style.<sup>15</sup> But then the agent's probability function will give the probability that an act has good consequences given that it was reached by rule following (or given that it is the Kantian act) a high value, and feeding this into the expected value equation will give the rule-satisfying act (or the Kantian act) the highest expected moral utility.

Similarly, we can acknowledge the often made point that sometimes consequentialist considerations support not going through consequentialist *deliberations*.<sup>16</sup> It may be that acting spontaneously in the situation at hand is known to have the best results—ducking, swerving, smiling, playing a drop shot, and the like are commonly best done straight off as the spirit moves one and without further ado. But in such cases the conditional probability of good results given that one acts without further ado will be high, or at any rate higher than the probability of good results given one acts after deliberation, and so, consequentialism will give the right result that one should in such cases act spontaneously. In such cases the consequentialist should hold that one ought to be *consequentially motivated* although one should not consciously reason consequentially.

It might be objected that this is an impossible position for the consequentialist to take up, given our account of consequentialism. How can the probability function give a high value to good consequences conditional on spontaneous action? For probability here means the agent's subjective probability function, that is, the agent's beliefs in quantitative guise; and the whole point of these examples of spontaneous action is that the agent acts without thinking. Accordingly, it might be suggested, he or she will not have any beliefs of the needed kind. The agent will not, for instance, *believe* that ducking the blow will have good results. There will not be enough time for that thought, only time to duck instinctively. Consider, however, the familiar example in the philosophy of perception where you drive past an advertising billboard without consciously registering what is on the billboard. You are later asked what the billboard was advertising, and to your surprise you are able to answer. This shows that you have seen the billboard and what was on it, although you were not conscious of the fact. In the same way one who ducks believes that not ducking will have results that are unpleasant despite the fact that the

15. An example of Railton's, "Alienation, Consequentialism, and the Demands of Morality," p. 116.

16. See, for three examples among many, Railton; Smart, p. 43; and the helpfully detailed account of a number of different examples in Pettit and Brennan.

thought to that effect is not consciously entertained. Spontaneous action is not action without belief, it is action without the conscious reviewing of belief.

What is central here is the distinction between the immediate springs of action and the processes that lead up to the obtaining of those springs. Otherwise consequentialism can seem to face a dilemma. Suppose that consequentialism says nothing about the mind of the agent at all. It says merely that right action is action with property  $\phi$ , for some consequentialist treatment of  $\phi$  which pertains solely to what in fact would happen and not at all to what the agent thinks. In that case, consequentialism, as Williams puts it, “has to vanish from making any distinctive mark in the world,” by which, I take it, he is, at least in part, making the point we made earlier that consequentialism must say something about right decision.<sup>17</sup> On the other hand suppose that consequentialism is expressed as a doctrine about how to go about making the morally right decision, as a variety of subjective consequentialism in Railton’s terms, and suppose in particular that it says to think along  $\phi$  lines. What then if thinking along  $\phi$  lines is discovered to have bad consequences in certain situations?<sup>18</sup> Our decision-theoretic account of consequentialism disarms the second horn of the dilemma by answering that in such situations the agent ought not to think along  $\phi$  lines, for the agent’s beliefs will then include that thinking along  $\phi$  lines in such situations has a low expected moral utility.

It is important to note that on our account, consequentialism is not committed to the view that maximizing expected moral utility is the right motive for action. A number of writers have made the point that doing something because you consider you ought to do it rather than doing it because you want to is, generally speaking, not the mark of the kind of person it is comfortable to be around. Being nice to someone solely because it is your duty to be nice to them is the kind of niceness we can all do without.<sup>19</sup> Michael Stocker sees this as a problem for what he calls “the standard view.” As he puts it, “The standard view has it that a morally good intention is an essential constituent of a morally good act. This seems correct enough. On that view, further, a morally good intention is an intention to do an act for the sake of its goodness or rightness.”<sup>20</sup>

17. Williams, “A Critique of Utilitarianism,” p. 135. Also his otherwise rather dark remarks on pp. 134–35 about utilitarianism retiring or being left with nothing more than “total assessment from the transcendental standpoint” may be a way of expressing the idea that utilitarianism must at some level be a species of decision theory. I am indebted here to a discussion with Thomas Scanlon in 1988.

18. A similar point can arise for deontological theories of course. “Keep your promises” is not in itself a rule of decision, though “Keep what you *take* to be your promises” is. But what if you are know that you are very bad at remembering what it is that you promised to do?

19. For some convincingly detailed examples showing this, see Michael Stocker, “The Schizophrenia of Modern Ethical Theories,” *Journal of Philosophy* 73 (1976): 455–66, and Railton.

20. Stocker, p. 462. Incidentally, on my view, consequentialism does not imply that a morally good intention is essential to a morally good act, at least if morally good act here

Perhaps this is true of some ethical views properly referred to as “standard,” but it is not true of consequentialism as characterized here. The right motive for action on the consequentialist view is the agent’s beliefs combined with desires that conform to the consequentialist value function, and the consequentialist value function *assigns no value as such to maximizing expected utility*. Take the drugs, mark 1 case. Prescribing drug A is the right thing to do because  $\sum_i Pr(O_i/A_j) \times V(O_i)$  takes a maximum value when  $A_j$  is prescribing drug A. But this fact that it takes a maximum value does not then confer *additional* value on prescribing drug A. That would be double-counting. What ought to move a person to action according to consequentialism are desires which may be represented as ranking states of affairs in the consequentialist way, but maximizing expected utility is not a factor in this ranking.

Before we turn to how our account of consequentialism helps with the nearest and dearest objection, I need to note an annoying complication. I have been arguing for an interpretation of consequentialism which makes what an agent ought to do the act which has the greatest expected moral utility, and so is a function of the consequentialist value function and the agent’s probability function at the time. But an agent’s probability function at the time of action may differ from her function at other times, and from the probability function of other persons at the same or other times. What happens if we substitute one of these other functions in place of the agent’s probability function at the time of action? The answer is that we get an annoying profusion of ‘oughts’. Consider the drug case, mark 1. I said that the intuitively correct answer to what Jill ought to do is prescribe drug A, and so it is. But suppose Jill later conducts a piece of definitive research which establishes that with patients of John’s blood type there is absolutely no chance that drug B will cause death, and in fact with such patients drug B is certain to effect a complete cure without any bad side effects. What will she then say about her past treatment of John, was it the right treatment or the wrong treatment? The natural thing to say would be something like, “By the light of what I now know, I ought to have prescribed drug B, but it would have been quite wrong to do so at the time.” But if it would have been the wrong thing to do at the time, how can it be what she ought to have done?

I think that we have no alternative but to recognize a whole range of oughts—what she ought to do by the light of her beliefs at the time of action, what she ought to do by the lights of what she later establishes (a retrospective ought, as it is sometimes put), what she ought to do by the lights of one or another onlooker who has different information on the subject, and, what is more, what she ought to do by God’s lights,

---

means what an agent ought to do. It is possible to do the right thing for the wrong reason. For an act which maximizes expected moral utility might also, as it happens, maximize expected highly immoral utility, and it might be that which prompts the agent to action. What is true is that doing an act for the right reason is sufficient but not necessary for it being what ought to be done *in the sense we are insisting is central in ethics*.

that is, by the lights of one who *knows* what will and would happen for each and every course of action.<sup>21</sup> The last will be a species of objective ought of the kind that features in Railton's (and Brink's) account of objective consequentialism. I hereby stipulate that what I mean from here on by 'ought,' and what I meant, and hope and expect you implicitly took me to mean when we were discussing the examples, was the ought most immediately relevant to action, the ought which I urged it to be the primary business of an ethical theory to deliver. When we act we must perforce use what is available to us at the time, not what may be available to us in the future or what is available to someone else, and least of all not what is available to a God-like being who knows everything about what would, will, and did happen.

It might be tempting to conclude that my acknowledgment that there are a variety of oughts means that I am not really disagreeing with one who urges that what a person ought to do according to consequentialism is what *in fact* has the best consequences; we are rather talking past each other. However, the substantive issue remains of the need for a moral theory to elucidate the ought most immediately relevant to action, and of how this should be done, quite independently of whether or not the target notion is unambiguously captured by 'ought' in English.

#### A REPLY TO THE NEAREST AND DEAREST OBJECTION

The decision-theoretic way of understanding consequentialism gives a major role to the agent's subjective probability function. This fact is the key to our reply to the nearest and dearest objection. I think that the reply can be most easily grasped by leading up to it via two examples: the drug example, mark 3, and the crowd control example.

##### *The Drug Example, Mark 3*

In mark 1, Jill had three drugs, A, B, and C, and one patient. This time Jill has three patients, A, B, and C, and one drug, and only enough of that drug to administer to a single patient. Her choice in mark 1 was between drugs, her choice this time is between patients, but it is all the same a similar style of choice situation that faces her. For we are given that she knows that patient A will derive considerable benefit from the drug without being completely cured, and also that one or other of patients B and C would be completely cured by the drug. However, she also knows that one or other of patients B and C would be killed by the drug. She has no way of telling which of B and C would be the one completely cured and which would be the one killed. What ought Jill to do?

The answer obviously is to administer the drug to patient A, and this is of course the answer our decision theoretic approach delivers.<sup>22</sup>

21. There are also the various nonmoral oughts—prudential etc., but that is another, and here irrelevant, dimension of variation.

22. Assuming of course that getting more information and then acting is not a viable option.

The expected moral value attached to administering the drug to A is higher than that for administering to B and higher than that for administering to C, because the possibility of a rather better result in those two cases goes along with a significant chance of a very much worse result. Of course, Jill knows that there is a better course of action open to her in the sense of a course of action which would have better consequences than administering the drug to A, but her problem is that she does not know whether it is administering the drug to B or administering it to C which is that better course.

What do we learn from this example to help us with the nearest and dearest objection? Well, it would clearly be a mistake to accuse Jill of an illegitimate bias toward patient A when she gave the drug to him rather than the others. Jill is biased toward patient A in the sense that her actions are directed toward securing his good, but the explanation for this fact is not that her preference function gives a greater weight to a benefit for A rather than one for B or for C. The explanation lies in her probability function. Consequentialism demands of us an impartial preference function, for its value function gives equal weight to the happiness, or preference satisfaction, or pleasure, or share of the ideal good, or . . . of each individual, but what the example tells us is that the fact in and of itself that our behavior is directed toward securing the happiness, or preference satisfaction or . . . of a small group—our family, friends, and so on—does not in itself show that we have an illegitimately biased preference function by the standards of consequentialism. The explanation of the directed nature of our behavior may lie in our probability functions.

The question, then, for consequentialists is the following. Can the special regard we have for a relatively very small group of people—to the extent that it is morally justified—be explained probabilistically, in terms of our special epistemological status with regard to our nearest and dearest, rather than in terms of an agent-relative preference function? The drug example mark 3 does not show that the answer to this question is yes. What it shows is that this is the key question that we need to ask.

I do not have a decisive argument that the answer to this key question is yes. What I do have are two considerations that suggest that it may well be yes. The first I will introduce with the crowd control example.

### *The Crowd Control Example*

Imagine that you are a police inspector who has been assigned the task of controlling a large crowd at a forthcoming soccer match. You have to choose between two plans: the scatter plan and the sector plan. The scatter plan is put to you in the following terms. “Each person in the crowd is of equal value. Any plan which told a member of the police squad to focus his or her attention on any particular person or group of persons would be immoral. Therefore, each member of the squad must roam through the crowd doing good wherever he or she can among as widely distributed a group of spectators as possible.” The sector plan is put to you in the following terms. “Each member of the squad should

be assigned their own sector of the crowd to be their special responsibility. This way members of the squad will not get in each other's way, and will build up a knowledge of what is happening in their sector and of potential trouble makers in it, which will help them decide on the best course of action should there be trouble. Also we will avoid a major problem for the scatter plan, namely, the possibility that at some particular time there will be part of the crowd which no one is covering. Of course, the sector plan should be administered flexibly. Although, as a general rule, each squad member should confine his or her attention to their assigned sector, if things are going particularly badly in another sector, and it is clear that an extra helping hand would make a big difference, then a transfer of attention may well be justified."

The plan that we follow in day to day life is, of course, the sector plan. We focus on a particular group, our family, friends, and immediate circle, while allowing that we may properly neglect them if the opportunity arises to make a very big difference for the better elsewhere. As the point is sometimes put, though it would be quite wrong to neglect family and friends in order to achieve a small increase in welfare elsewhere, it would be quite proper to neglect them in order to achieve peace in the Middle East. And we can approach the question of whether probabilistic considerations can provide a justification from the consequentialist's point of view for our focusing on family and friends, by asking when the sector plan would be the right plan for a consequentialistically minded inspector to adopt.

The sector plan would be strongly indicated in the following circumstances. (a) When getting to know certain individuals is important for achieving good results. The scatter plan distributes any given squad member's attention very widely, making any detailed knowledge of the psychology of particular individuals difficult. If good results depend on such knowledge, then squad members should restrict themselves to a smaller group, as in the sector plan. (b) When achieving good results involves coordinating a series of actions. Sometimes an isolated action has little effect in itself. What is needed is an extended plan of action, with later actions chosen on the basis of positive and negative feedback from the results of earlier actions. Think of the contrast between a one shot drug treatment and an extended course of treatment with later drugs and dosages being chosen in the light of the effects of earlier treatments. (c) When achieving good results depends on setting up mutual trust and respect and understanding between individuals. The traditional "bobby on the beat" is a special kind of sector arrangement and rests on exactly this kind of point. (d) When there is a significant chance of different squad member's actions nullifying each other if directed toward the same people. When we are in a situation where "too many cooks spoil the broth," the sector plan is clearly superior to the scatter plan. (e) When there is an obvious way to assign police to separate sectors which coincides with their natural inclinations and enthusiasms, particularly

when this fact is common knowledge. This reduces, and is known to reduce, the setting up costs of the sector plan by avoiding costly debate over who takes responsibility for which sector. It also increases the penalty consequent upon a squad member not policing their natural sector, by increasing the chance that that sector will remain unattended through other police wrongly assuming that it is attended.

Clearly, there is a great deal more to be said here, much of it to do with straight empirical facts.<sup>23</sup> But I hope that I have said enough to make it plausible that the sector plan is indicated in the kind of circumstances that apply in our day-to-day interactions with the world around us and the people in it. It is hard to know what actions will have good effects, and our opinions on the matter are much better founded in the case of people we know well precisely because we know them well. Achieving good results is very often a matter of coordinating a series of actions rather than scattering largesse around. Mutual trust and affection are important for good results. Too many cooks can spoil the broth when it comes to interacting in a beneficial way with one's fellow human beings. There is very obviously a group of people whose welfare we are naturally inclined to concern ourselves with, namely, those nearest and dearest to us. And, most important for our decision-theoretic approach to consequentialism, facts such as those just adumbrated are, I take it, pretty much common knowledge.

My suggestion, then, is that the consequentialist can reply to the nearest and dearest objection by arguing that the kind of direction of attention toward those we hold dear which is so characteristic of a worthwhile life can be explained without attributing a biased value function. It is instead a reflection of the nature of our probability functions, in particular, of the kinds of facts about the epistemology of achieving good consequences that we have been rehearsing. The suggestion is not of course that the kind of direction of attention we typically manifest *in fact* toward those we hold dear can be explained without attributing a biased value function. It is no objection to consequentialism that, according to it, we ought to do more than we in fact do for people we hardly know. We ought to do more for people we hardly know. We are too tribal. The suggestion is that a considerable degree of focus on our family and friends, enough to meet the demand that our lives have a meaningful focus, is plausibly consistent with living morally defensible lives according to consequentialism.

### *On Three Objections*

1. Williams has argued that "it [consequentialism] essentially involves the notion of *negative responsibility*: that if I am ever responsible for anything,

23. See, e.g., the discussion of the allocation of responsibilities in Philip Pettit and Robert Goodin, "The Possibility of Special Duties," *Canadian Journal of Philosophy* 16 (1986): 651–76.



then I must be just as much responsible for things that I allow or fail to prevent, as I am for things that I myself . . . bring about. Those things must also enter my deliberations, as a responsible moral agent, on the same footing. . . . What matters [according to consequentialism] with respect to a given action is what comes about if it is done, and what comes about if it is not done, and those are questions not intrinsically affected by the nature of the causal linkage, in particular by whether the outcome is partly produced by other agents."<sup>24</sup>

If Williams is right, we are in trouble. The key idea behind our reply to the nearest and dearest objection was that the reflections we grouped under the heading of the sector plan made it plausible that a consequentialist ought to take special responsibility for what is within his or her ken, and that will obviously involve making *who* does something a very important matter in many cases—and that runs directly counter to Williams's claim that who does something is irrelevant for consequentialists. However, it is crucial here to bear in mind the distinction between value and expected value. Williams is right that consequentialism's value function gives no weight per se to who does something (and that no doubt was what he had in mind), but nevertheless who does something can be enormously important to the expected value of a course of action, and it is that which is crucial according to our account of consequentialism, and that is how who does something can "enter my deliberations." In particular, Smith may hold that it would, all in all, be better were A to obtain than were B to obtain, but it does not follow that Smith qua consequentialist should seek to bring about A rather than B. For Smith may hold, in addition, that someone else, Jones, knows more about the matter than he himself does and that Jones has good values. In such a case, the decision facing Smith is between he himself bringing about A rather than B, or instead leaving the decision to Jones as to whether or not to bring about A rather than B, and it is easy to show that the latter may have the greater expected moral utility for Smith. The crucial point is that, though for Smith the probability of good consequences given he does B is low, the probability of good consequences given *Jones* does B may be high because of Smith's opinion that Jones is best placed to make the decision. In general, in cases where we judge it best to leave a decision between A and B to the experts, as we say, although we may have ourselves a view as to which of doing A and doing B has the greatest expected value, leaving the matter to the experts may have the greatest expected value of all. Thus, *who* does something *can* be crucial according to consequentialism.

It might be replied on Williams's behalf that he did not have expected value in mind and that the point we have just made only holds for expected value. However, this would make nonsense of Williams's (correct)

24. Bernard Williams, "Consequentialism and Integrity," in Scheffler, ed., p. 31. Williams is, of course, supposing that *effects* due to the identity of the agent have been incorporated into the consequences.

insistence that consequentialism be an ethical *decision* theory and of the talk in the above quotation of how things “must enter my *deliberations*” (my emphasis).

2. Railton raises (as a preliminary to a reply), the nearest and dearest difficulty for consequentialism with the following example. “Juan and Linda . . . have a commuting marriage. They normally get together only every other week, but one week she seems a bit depressed and harried, and so he decides to take an extra trip in order to be with her. If he did not travel, he would save a fairly large sum that he could send Oxfam to dig a well in a drought-stricken village. Even reckoning in Linda’s uninterrupted malaise, Juan’s guilt, and any ill effects on their relationship, it may be that for Juan to contribute the fare to Oxfam would produce better consequences overall than the unscheduled trip.”<sup>25</sup> It might be objected that what I have said so far in no way meets the objection posed by this example. But from the decision-theoretic point of view what is crucial is not that “it *may* be that for Juan to contribute the fare to Oxfam would produce better consequences,” it is how likely it is to do so. And, of course, the effects of isolated acts of charity on the Third World are a matter of considerable debate, whereas Juan can be pretty certain of at least some of the effects of his making the unscheduled trip. It is important here to remember that the relevant consequence of sending, say, \$500, should not be approached by asking what \$500 will buy in the Third World, but by addressing the likely *differences* between what would we achieved by the sum Oxfam would have without Juan’s \$500, and what would be achieved by the sum with Juan’s \$500.<sup>26</sup>

What plausibly is obvious is that many of us in advanced Western societies could achieve a great deal more good if we devoted our energies to a systematic, informed program of transferring any excess wealth toward the Third World. I do not mean isolated donations of airfares, but neither do I mean just sending a lot more money until it really hurts. I mean becoming actively involved in and knowledgeable about what is going on in the Third World: learning how aid agencies work, which ones do good, which, knowingly or unknowingly, do harm; finding out exactly how villages use money sent to them; the effects outside money and services typically have on the local social and economic structures; and so on and so forth. But how can this observation possibly constitute a nearest and dearest objection to consequentialism? A person who behaved in the way that I have just described *would* be directing her attention to those nearest and dearest to her. For she would be paying special attention to a relatively small section of the world’s population, and she would be

25. Railton, p. 120.

26. You may not like this way of approaching the consequences to be assigned to giving \$500 to Oxfam, perhaps influenced by the examples in Derek Parfit, *Reasons and Persons* (Oxford: Oxford University Press, 1984), chap. 3, but that is a *different* objection—and in my view a tempting but mistaken one (see Frank Jackson, “Group Morality,” in Philip Pettit et al., *Metaphysics and Morality* [Oxford: Blackwell, 1987], pp. 91–110).

giving a special place to her own projects, one of which would precisely be helping certain people in the Third World. For her, the people whose welfare she was particularly concerned with would be those people she had studied, and so got to know and understand, living in various villages in the Third World, as opposed to those living in the same house or the same neighborhood as she does.

3. It might well be objected that we can distinguish *two* nearest and dearest objections, and that I have replied to only one of them. One objection is, "How can consequentialists make sense of the fact that there is a relatively small group of people whose welfare plays a special role in our lives, given the agent-neutral nature of consequentialism's value function?" Our reply was that consequentialism should be viewed decision theoretically. The way right value translates into right action is through an agent's beliefs, and that when this is appreciated, empirical facts about our cognitive powers and situation make it plausible that our actions should be highly focused much of the time. The other objection is, "How can consequentialists make sense of it being the *particular* small group of people that it mostly is?" Perhaps consequentialism can make sense of there being a small group, but why the small group of family, friends, fellow citizens, and the like that it so often is?

One possible reply is that consequentialism cannot make sense of this but that that is no objection to it. We *are* outrageously tribal in our everyday morality and, hence, so much the better for consequentialism that it makes this clear. I cannot believe this. I grant that we are unduly tribal but not that we are outrageously so. I think that we can give a consequentialist explanation of why, for most of us, the special group is our family and friends in terms of empirical facts about human character and psychology.

One way you might draw on empirical facts about human nature is to argue that some particular action giving preference to family and friends in a way which goes against consequentialist principles is wrong but excusable in some sense because it is the exercising of a character which is good in consequentialist terms. This is William Godwin's claim about his famous example of your having to choose between rescuing Fenelon, a famous author and archbishop, and a valet who happens to be your own father, from a burning house: rescuing your father is the wrong action but at the same time the action which springs from the right character.<sup>27</sup> The idea is that although there is perhaps in theory a better character which would lead to the best action as judged consequentially, in practice such a character is not available to us, or at least not to most of us.

27. William Godwin, *Enquiry concerning Political Justice* (Oxford: Oxford University Press, 1971). The example is given on p. 71; the point about character was something of an afterthought prompted by the reception that greeted his answer that you ought to abandon your father (see p. 325). Railton takes a similar position on his example quoted earlier about a commuting marriage.

I think that there is an element of “ducking the question” about this reply. If the action which favors family and friends is right and consequentialism says that it is wrong, then consequentialism is false, and there’s an end on it. If, on the other hand, the claim is that it is the action which favors family and friends which is wrong, then it is *that* which needs to be established, not facts about good character. Or perhaps what we are being offered is a *variant* on consequentialism according to which an action is to be judged not directly but via the status, judged consequentially, of the character which gives rise to it, but then we appear to be landed with a dubious compromise reminiscent of rule utilitarianism.<sup>28</sup> If consequences are the key in one place, why not across the board?

I am not here denying the correct and important point that some particular action may be wrong in consequentialist terms and yet spring from a character which is right in consequentialist terms.<sup>29</sup> I am denying that the point helps with the essentials of the nearest and dearest objection. For the consequences of having a character which gives a special place in one’s affections and concerns to those persons who are closest to one are, in the main, consequences of the *manifestations* of such a character, that is, of the actions which are especially directed to the needs of those closest to one. Hence, a consequentialist justification of such a character presupposes a consequentialist justification of those actions—which returns us to the very question raised by the nearest and dearest objection. Of course it is not the case for every character trait that the consequences of possessing it are, in the main, the consequences of manifesting it. A major consequence of possessing a certain character trait may be that people know that you possess it, which knowledge may in turn have a major effect on their behavior without it ever being necessary to manifest that disposition. Being disposed to react with pointless violence on being attacked is an example; a nation’s disposition to make a nuclear response to a major nuclear attack is another possible example familiar in the literature on nuclear deterrence. Our point, though, is that the character trait of being especially concerned with the welfare of those closest to us is, like most dispositions, known by and large through its manifestations and has its effects principally via those manifestations. I know that you are especially concerned about the welfare of your family because your actions display that concern. Hence, a consequentialist justification of that character trait awaits a consequentialist justification of those actions.<sup>30</sup>

28. Railton makes it clear that he is not offering such a variant. I am not so sure about Pettit and Brennan.

29. See, e.g., Railton; Sidgwick; and Parfit, sec. 14.

30. Nor am I denying the relevance of the distinction between evaluation of action and evaluation of character in explaining the very mixed feelings we would have toward a father who saved a stranger in preference to his daughter on the ground that he happened to know that the stranger was slightly more worthy of saving from an agent-neutral point of view. The distinction enables the consequentialist to explain the mixed feelings as a response to witnessing simultaneously the manifesting of a wrong character in a right action. I think that this is part of the purpose to which Godwin wished to put the distinction.

Be all this as it may, I think that points about character and human nature can be put to more direct work here. One's character can be a major factor in settling what consequences are likely, and so can be a major factor in settling what acts are right from the (act) consequentialist's point of view.

Some actions are such that they only have good results if they are followed up in the right way. Taking the first capsule in a course of antibiotics will only have good results if the remainder are taken at the right times; agreeing to write a book review will only have good results if you write the review in good time; going on a beach holiday will only have good results if you avoid getting badly sunburnt; and so on and so forth. In all these cases it is better not to start if you are not going to follow up in an appropriate way. From the consequentialist perspective whether action A ought to be done depends in part on what the agent would in fact do subsequent to doing A.<sup>31</sup>

This means that in deciding what to do here and now an agent must take account of what he or she will do in the future, and that involves taking very seriously questions of character. Do I have the persistence that will be called for, will I remain sufficiently enthusiastic about the project to put in the time required, will I be able to retain a sufficiently impartial outlook, will I be able to avoid the various temptations that will arise, and so on and so forth? For some of us in some situations these kinds of considerations count *against* attempting to secure benefits for our friends and family. We do better sometimes with people we are not so close to. Some men should most definitely not play doubles in tennis with their wives as partners. But as a rule we do better for reasons of character (that no doubt have an evolutionary explanation) with projects that involve family and friends rather than strangers. This is simply because we are much less likely to lose the enthusiasm required to see the project through to a successful conclusion when the project benefits people we have a particular affection for. Perhaps a mundane example will make the point clearer. Jones may be able, in principle, to do an equally good job of organizing the seminars in history or the seminars in philosophy for the forthcoming year. She has the required knowledge and contacts in both areas. She is however much more excited by philosophy

31. In Frank Jackson and Robert Pargetter, "Oughts, Options, and Actualism," *Philosophical Review* 95 (1986): 233–55, we argue that this is true in general, not just for consequentialism, provided only that consequences sometimes play a central role in determining what ought to be done. This view is controversial; see the extensive literature referred to therein. In Frank Jackson, "Understanding the Logic of Obligation," *Proceedings of the Aristotelian Society* 62, suppl. (1988): 255–70, I argue that the best way to understand what is going on is in terms of early actions of the agent being actions of early temporal parts. (In both papers the argument is conducted, for reasons of expository convenience, mainly in terms of what *objectively* ought to be done, rather than the decision theoretic ought that I have put in center stage here. I now think that this was an unfortunate choice of argumentative strategy: it obscured the point I have highlighted here—that ethics pertains most particularly to acting.)

than by history. In such a case, even prescind from her own enjoyment, it may well be that she ought to take on the task of organizing the philosophy seminars. For, although she knows that she *could* do equally well at either, she knows that she *would* most likely do better if she takes on the philosophy program, and her knowledge of this fact should influence her to agree to take on the philosophy program rather than the history program.

We have seen that the good consequentialist should focus her attentions on securing the well-being of a relatively small number of people, herself included, not because she rates their welfare more highly than the welfare of others but because she is in a better position to secure their welfare. Typically, this will involve her in settling on a relatively extended program of action which will take some resolution and strength of character to carry forward successfully. Before she starts she knows, if she is at all like most of us, that the chances of success are much greater if she makes the relatively small group those who are her family and friends, rather than those she hardly knows. There are exceptions to this generalization about human psychology, perhaps Mother Teresa is one, perhaps Ralph Nader is another; from reports it seems that they have the ability to carry through a demanding program of action which benefits a group of people which, though tiny by comparison with the population of the world, is large by comparison with the circle of family, friends, and associates that provide the principle focus of action for most of us. They do not seem to be dependent on the kind of close personal relationships that are essential to keep most of us from being outrageously selfish.

## CONCLUSION

Consequentialism tackles the question of what an agent ought to do in terms of the values of outcomes and assigns those values in an agent independent way, and yet the lives we consider worth living give a quite central place to certain of our fellow human beings. We have an agent relative moral outlook. My argument has been that the consequentialist can plausibly explain agent relativity in terms of the role probability plays in the recovery of what an agent ought to do from the consequentialist's value function. The injunction to maximize *expected* moral utility, when combined with the facts we listed under the heading of the sector plan, means that the consequentialist can accommodate our conviction that a morally good life gives a special place to responsibilities toward a smallish group. *Which* smallish group is another question, and here I argued that for most of us the group should be chosen tribally. Because of empirical facts about our natures, that choice decreases the chance that we will backslide.

One objection to consequentialism is that it conflicts with firmly held moral convictions, in particular concerning our obligations toward our nearest and dearest. It may be urged that my reply to this objection is seriously incomplete. For we can reasonably easily describe a possible

case where the factors I mentioned as providing a justification in consequentialist terms for favoring one's nearest and dearest do not apply, and yet, according to commonsense morality, one should favor, or at the least it is permissible to favor, one's nearest and dearest.<sup>32</sup> My concern, though, has been to reply to the objection that consequentialism would, given the way things more or less are, render the morally good life not worth living. I take this to be the really disturbing aspect of the nearest and dearest objection to consequentialism. Consequentialists can perhaps live with the conflict with commonsense morality, drawing for instance on the notorious difficulties attending giving a rationale for its central features.<sup>33</sup> But it seems to me that they cannot live with the conflict with a life worth living, given the way things more or less are. That would be to invite the challenge that their conception of what ought to be done had lost touch with *human* morality.

32. I am indebted to David Lewis and Kim Sterelny for forcibly reminding me of this fact.

33. See, e.g., Shelly Kagan, *The Limits of Morality* (Oxford: Oxford University Press, 1989).