

MORALITY, RULES, AND CONSEQUENCES

A Critical Reader

Edited by Brad Hooker,
Elinor Mason, and Dale E. Miller

EDINBURGH UNIVERSITY PRESS

(2000)

utilitarianism ranks particular acts. When acts alone are ranked, however, all other conditions are given. It is unclear how to rank all of these particular items when the other conditions that affect a particular item's utility include particular items in other classes whose members are likewise to be ranked and for that reason would not seem to be regarded as given.

19. Railton, 'How Thinking about Character', p. 41f.
20. Note that my comments on this example are not purely utilitarian. For the aim of valoric utilitarianism is not to maximise utility but to use generic-utilitarian elements in the construction of a defensible criterion of right conduct. I'll get back to this point in a moment.
21. Railton, 'How Thinking about Character', p. 410.
22. Rawls, *A Theory of Justice* (Cambridge, MA: Harvard University Press, 1971), p. 48.
23. Amartya Sen, 'Utilitarianism and Welfarism', *Journal of Philosophy*, 76 (1979), pp. 463–89, at 463.
24. *Ibid.*, p. 464.
25. *Ibid.*, p. 466. Sen's 'utilitarian moral structures' support moral appraisals of acts only, although outcome utilitarianism would not seem to imply any such restriction and utilitarian moral appraisals would seem possible of things other than acts.
26. *Ibid.*, p. 464.
27. *Ibid.*, pp. 463–89, at p. 466. Although Sen refers to an argument establishing this point (p. 466, n. 8), I have not found it.
28. *Ibid.*, p. 466.
29. We could avoid this problem by applying the influencing variables and balancing them in an 'intuitionistic' manner. This would not yield a determinate principle, but it suggests a model for utilitarian reasoning in practice.
30. I think it is also unclear how to determine whether a principle best or most faithfully expresses the guiding utilitarian idea. Call that the representation problem. I'll come back to it in the next section.
31. Here's a further complication. The range of principles that we imagine different individuals conforming to might be limited to ones that are recognisably utilitarian; alternatively, we might include within the range of principles some that are not recognisably utilitarian.
32. It might be objected that morality requires its principles to apply universally. The alternative criterion does not deny this. It accords with the possibility of fundamental moral disagreement, which the concept of morality does not deny.
33. The previous complication might apply here. The range of principles that we imagine individuals conforming to might either be limited to ones that are recognisably utilitarian or might include some that are not recognisably utilitarian.
34. It might be objected that morality requires its principles to apply constantly. The alternative criterion does not deny this. It accords with the possibility that individuals change their basic moral commitments over time, which the concept of morality does not deny.
35. Because acceptance involves certain moral attitudes, whereas simple conformity does not, acceptance may also involve further distinguishing consequences.
36. This alternative test would also seem to reflect a more realistic notion of the role we might ideally expect moral principles to play in our lives than perfect conformity does.
37. It is also conceivable that some neutral factor, such as the concept of morality or of rationality, would, when added to generic utilitarianism, somehow disqualify all but one potential moral representative of the guiding utilitarian idea and would thus save utilitarianism from moral ambiguity. But I see no reason to assume this.

Global Consequentialism

Philip Pettit and Michael Smith

SECTION I: GLOBAL VS LOCAL CONSEQUENTIALISM

Global consequentialism identifies the right x , for any x in the category of evaluands – be the evaluands acts, motives, rules, or whatever – as the best x , where the best x , in turn, is that which maximises value. Value may be actual or expected, of course, but we shall ignore that complication here. So, for example, according to global consequentialism, the right act for someone to perform is the act that has greater value than any of the acts that might have been performed instead; the right motive-set for someone to have is that motive-set whose possession has greater value than any of the motive-sets that might have been possessed instead; the right set of rules for someone to have internalised is that set of rules which has greater value than any of the sets of rules that they might have internalised instead, and so on.¹

The best-known statement and defence of global consequentialism is to be found in Derek Parfit's *Reasons and Persons*. Here is what Parfit has to say:

There are different versions of *Consequentialism*, or *C*. *C*'s central claim is

(C1) There is one ultimate moral aim: that outcomes be as good as possible
C applies to everything. Applied to acts, *C* claims . . .

(C2) What each of us ought to do is whatever would make the outcome
 best

. . . Consequentialism covers, not just acts and outcomes, but also desires, dispositions, beliefs, emotions, the color of our eyes, the climate, and everything else. More exactly, *C* covers anything that could make outcomes better or worse. According to *C*, the best possible climate is the one that would

make outcomes best. I shall use 'motives' to cover both desires and dispositions. C claims

(C5) The best possible motives are those of which it is true that if we have them, the outcome will be best.²

Thus, according to Parfit, C tells us not just which acts are the right ones to perform, but also which desires, beliefs, and emotions are the right ones to have, whether it is right to have this colour eyes or that, whether it is right for it to be rainy or cloudy or sunny, and so on and so forth. This is global consequentialism if anything is.³

As Parfit's remarks make clear, the crucial feature of global consequentialism is that it does not privilege any category of evaluand. In particular, it does not privilege the category of acts that has often been privileged, by default, in much consequentialist writing. It does not say, for example, that the right motive-set for someone to have, or the right set of rules for someone to have internalised, is that set which would promote the choice of the right acts. This is important, as someone's possession of certain motives, or his or her having internalised certain rules, may have consequences that are not mediated by any act to which those motives or rules give rise.⁴ Your clear benevolence towards me, and mine towards you, can provide each of us with a sense of warmth and reassurance independently of any acts that it occasions. And the mere knowledge that you have internalised a rule of promise-keeping provides me, well in advance of any contract we enter into, with a rich sense of the arrangements we may form.

Global consequentialism thus contrasts with any form of consequentialism that privileges one or another category of evaluand, and which is therefore *local*. The sort of consequentialism just described – the sort that privileges acts – is a local act consequentialism. It privileges the category of acts, defining right acts as those which maximise value – in this it agrees with global consequentialism – but it then goes on to identify right motive-sets, right sets of rules, and so on, as the sets that promote the choice of right acts. Other forms of local consequentialism are distinguished from global in so far as they privilege some non-act category of evaluand.

Consider, for example, the version of rule utilitarianism described by J. J. C. Smart:

Rule utilitarianism is the view that the rightness or wrongness of an action is to be judged by the goodness or the badness of the consequences of a rule that everyone should perform the action in like circumstances.⁵

This form of rule utilitarianism privileges rules, and then defines right acts in terms of the right rules. It is therefore local, in the sense that contrasts with global.⁶

The version of rule consequentialism Brad Hooker defends is similarly local. According to Hooker,

An act is wrong if and only if it is forbidden by the code of rules whose internalization by the overwhelming majority of everyone everywhere in new generations has the highest expected value. The evaluation of a code is to count all the costs involved in getting it internalized.⁷

Hooker's version of rule utilitarianism thus also privileges rules, and defines right acts in terms of the right rules.

Robert Merrihew Adams recommends that motive consequentialists embrace a form of local consequentialism.

This position – that we have a moral duty to do an act, if and only if it would be demanded of us by the most useful kind of conscience we could have – may be called 'conscience utilitarianism', and is a very natural position for a motive utilitarian to take in the ethics of actions.⁸

Conscience utilitarianism thus privileges consciences, and defines right acts in terms of the conscience it would be best for us to have.

Finally, and more generally, the various forms of indirect utilitarianism described by Simon Blackburn are all local theories that privilege non-act evaluands.

The doctrine that applies utilitarianism to actions directly, so that an individual action is right if it increases happiness more than any alternative, is known as direct or act utilitarianism. Indirect versions apply in the first place to such things as institutions, systems of rules of conduct, or human characters: these are best if they maximize happiness, and actions are judged only in so far as they are those ordained by the institutions or systems of rules, or are those that would be performed by the person of optimal character.⁹

Thus, according to Blackburn, indirect utilitarianisms quite generally privilege some non-act category of evaluand – institutions, rules of conduct, human characters, or whatever – and then go on to define right actions in terms of the right things in the privileged category.

The form of local consequentialism that privileges acts has been justly criticised in the literature by consequentialists themselves.¹⁰ Their objection has been that by ignoring the non-act mediated consequences of motives, rules, and the like, local-act consequentialism fails to pay due regard to the fact that these things are the significant sources of value that they so plainly are. Strikingly, however, consequentialists have said little or nothing about the corresponding failings of the various forms of local consequentialism that privilege some non-act category of

evaluand. Indeed, as we have seen, some have even gone on to defend their own versions of local consequentialism.¹¹

In our view, this oversight is unjustified. Consequentialists should think that those forms of local consequentialism that privilege some non-act category of evaluand, and redefine rightness in such terms, are just as bad as those that privilege acts. These forms of local consequentialism also fail to pay due regard to a significant source of value. In this case, however, the significant source of value is *acts*.¹²

SECTION 2: A PROBLEM AND OUR STRATEGY FOR OVERCOMING IT

There is a problem to be overcome in mounting an argument against local consequentialism, however. What, after all, does it mean to privilege a particular category of evaluand in the definition of right acts? Unfortunately, because different answers can be given to this question, local consequentialism is difficult to formulate in a definitive way. The upshot is thus that even if we could provide a decisive refutation of some version of the theory, a committed local consequentialist could with some justification insist that the refuted theory isn't the best version. We will illustrate this problem with two examples.

Consider local motive consequentialism. According to this theory, the right motive-set for someone to have is that motive-set whose possession has greater value than any of the motive-sets that might have been possessed instead. In this respect local motive consequentialism agrees with global consequentialism. But how does local motive consequentialism go on to privilege the category of motives in defining right acts? There are various possibilities. It might go on to define right acts to be those which are caused by the right motives; or it might define them to be those which would have been caused by possession of the right motives; or it might define them to be those which would have been caused by the motives that it would be best for someone to try to have; or it might define right acts in some yet further way.

Or consider local rule consequentialism. According to this theory, the right rules for someone to have internalised is that set of rules adherence to which has greater value than any of the sets of rules that the person might have internalised instead. Again, in this respect it accords with global consequentialism. But how does local rule consequentialism then go on to define the right acts? It might define them to be those which are caused by having internalised the right rules; or it might define them to be those which would have been performed by someone who had internalised the right rules; or it might define them to be those which would have been performed by someone who had internalised the rules that it would be best for someone to try to internalise; or it might define right acts in some yet further way. Again, there is a variety of possibilities.

Our strategy, in order to overcome this problem, is to argue inductively against all versions of local consequentialism at once. We begin by considering the various versions of local motive consequentialism just described. We note that they are liable to a similar pattern of objection. In our view, every form of local consequentialism will be liable to that same pattern of objection, and we attempt to prove that this is so by indicating the way in which the objections apply to the various versions of local rule consequentialism just described. Once one sees how the pattern of objection applies in one case, we think it is hard not to conclude that all forms of local consequentialism will fall to that same pattern of objection. As with any inductive argument we may, of course, be wrong. However we leave it to the defenders of local consequentialism to prove that that is so.

SECTION 3: LOCAL MOTIVE CONSEQUENTIALISM: FIRST VERSION

Consider the version of local motive consequentialism which holds that the right acts are those which are caused by the right motives. We have two related objections to this theory. The first is that it is absurd to hold, as this theory does, that only acts which are caused by the right motives are right. The second is that it is equally absurd to hold, as this theory also does, that every act which is caused by the right motives is right.

Why hold that only acts which are caused by the right motives are right? After all, those people who do not have the right motives – and hence those whose acts cannot be caused by the right motives – are still able to act in ways that contribute more or less value to the world. For example, they are able to perform, or not to perform, those acts, of the acts available to them, that produce the best consequences, and they are also presumably able to perform, or not to perform, those acts, of the acts available to them, that would have been performed by people who did have the right motives. Given that the acts available to these people contribute different amounts of value to the world it surely follows that some of them are to be recommended over others, on consequentialist grounds. But once this point is conceded we must ask what merit there is in insisting, as the first version of local consequentialism insists, that the act thus recommended is still not the right act for these people to perform. It seems unnecessarily ad hoc to invent a new word to characterise the acts which are to be recommended to those who do not have the right motives.

The second objection is that it is equally absurd to suppose that every act which is caused by the right motives is a right act. Focus on a particular case in order to see why. Imagine that, in a certain situation, the best consequences will flow from someone's having a set of motives that includes a strong desire to provide benefits, no matter how small, for his children, and only a weak motive to provide benefits, no matter how large, for strangers. Moreover, imagine that this person has these

motives, and that he faces a particular choice situation in which he can provide either a smaller benefit for his children or a larger benefit for some strangers. Finally, let's suppose that in these particular circumstances, providing the larger benefit for the strangers will have better consequences than giving the smaller benefit to his children. Which is the right act for him to perform?

According to the first version of local motive consequentialism the right act for this person to perform is giving the smaller benefit to his children. But why? Let's agree that the person we are imagining *will* act so as to give the smaller benefit to his children, provided he functions properly psychologically. So much follows from the fact that his desire to give them the smaller benefit is stronger than his desire to give the larger benefit to the strangers. But the fact that he *will not* give the larger benefit to the strangers is neither here nor there in deciding whether his giving the strangers that larger benefit is the right thing for him to do.

Nor is the fact that, given that he has the motives that it is right for him to have and functions properly psychologically, the act of giving the larger benefit to the strangers is only available to him by grace of the possibility of a psychological malfunction or anomaly. The value or disvalue associated with psychological malfunction is, after all, to be decided by a consequentialist on consequentialist grounds, and, in the circumstances that we are imagining, giving the larger benefit to the strangers has the best consequences, notwithstanding the fact that it would involve psychological malfunction. A consequentialist should surely think that there is therefore much to rejoice about, and little to regret, if the person we are imagining manages to give the larger benefit to the strangers.

The latter point needs emphasising. We think that one main attraction of the view that it is always right to act on the motives that it is right to have is the implicit assumption that failing to act on these motives will somehow undermine possession of them. But this thought is multiply confused. For one thing, it is implausible to suppose that the occasional occurrence of a psychological malfunction which causes someone to act contrary to the motivations that he has will undermine his possession of these motives. Patterns of motivation are more robust than that. For another, if, as we are supposing, giving the larger benefit to the strangers has the best consequences then, even if it does undermine the agent's possession of the motives that it is currently right for him to have, by hypothesis the benefits of his so acting clearly outweigh that cost. It is thus hard to see why a consequentialist should be worried by the fact that the agent undermines his possession of the motives that it is currently right for him to have if, as we have supposed, the possible world in which he no longer has these motives contains more value than the world in which he retains them.¹³

But, in that case, we must surely ask what reason there is for supposing that it is always right for a person to act on the motives that it is right for him to have. The

only defence of the suggestion left, as far as we can see, is that the acts which are deemed to be right, on such a definition of rightness, accord more closely with our pre-theoretical intuitions about which acts are right and which are wrong than do the acts which are deemed to be right by a direct application of the consequentialist principle itself. But it seems to us that someone could think this only if she forgets about a whole range of possible cases in which people have and act upon the motives that it is right for them to have.

Consider, for example, the possible world in which there is a mad scientist who will make millions of people miserable if certain individuals don't have malignant motives, but who couldn't care less which acts these individuals perform. In such a world consequentialists will very plausibly suppose that it is right for these individuals to have malignant motives. But does it accord well with our pre-theoretical views about which acts are right and wrong to suppose that their acting on these malignant motives is right in that world? Not at all. If there is an act available to these individuals that produces greater value than that produced by acting on their malignant motives, then it surely accords much better with our pre-theoretical intuitions to suppose that the right act for them to perform in that possible world is the one that produces greater value.

The upshot, in our view, is thus that though consequentialists are right to concern themselves with the intuitive plausibility of their theory, the intuitions that should matter to them are the basic intuitions about what is valuable and what isn't. Those who find themselves reluctant to believe that right acts are those which maximise value should therefore reassess their entire commitment to a consequentialist way of thinking, rather than attempt to redefine rightness in the manner recommended by the first version of local motive consequentialism. Such a redefinition can be given no clear rationale.

SECTION 4: LOCAL MOTIVE CONSEQUENTIALISM: SECOND VERSION

Consider next the version of local motive consequentialism which holds that right acts are those which would have been caused by possession of the right motives, rather than those that are actually so caused. This version of the theory differs from the first version because it explicitly allows that people who do not have the right motives might nonetheless act rightly.

Thus, if we once again imagine that the right set of motives for people to have includes a stronger desire to give a small benefit to their children and a weaker desire to give a large benefit to strangers, then the main difference with the first version of local motive consequentialism is that, according to this second version, those who don't have this pattern of motivation are still able to act rightly just so long as they are able to give their children that smaller benefit. Of course, given that they have a much stronger desire to give a larger benefit to strangers it follows that they could

only do so if they were to act out of an even stronger desire to be a decent parent, or to avoid the criticism of other family members, or perhaps by acting on the basis of some psychological malfunction or anomaly. But, whatever the cause of their so acting, so long as they succeed in giving their children that smaller benefit they do at least succeed in acting rightly, according to this second version of local motive consequentialism. The real question facing this version, however, is the same as the question that arose in our discussion of the first version. What is the consequentialist case supposed to be for the claim that someone who acts as if they had the right motives always acts rightly?

Note, to begin, that since the person who acts as if they had the right motives may not actually have the right motives, there need not be the same non-act mediated benefits associated with their action as would have been in place if they did have the right motives. There may be some overlap – pretending to be a friend might have some of the good consequences of being a friend – but this is bound to be only partial and, in any case, highly contingent. There is thus little room on this front for making a consequentialist case for someone's acting as if they had the right motives.

Perhaps the thought is that acting as if you had the right motives will generally produce the right motives; that therein lies the motive consequentialist case for so acting. But this thought is confused. There are three points to be made against it. The first is that acting as if you had the right motives might not bring it about that you have the right motives. Indeed, the right motives might not be motives that you could bring about by any act available to you. Perhaps the only way to have acquired these motives is by having had a different childhood. The second is that if the right act is the act that the right motives would cause, then it is in any case not clear why it is right for anyone to try to bring about the right motives. After all, the act of trying to bring about the right motives is not one that the right motives would cause; the presence of the right motives would make the act of trying to bring about the right motives unnecessary. And the third point is that, among the motives that you could bring about in yourself, the motives that would be brought about by acting as if you had the right motives might be manifestly inferior to some other set of motives that you could have brought about by acting differently. Perhaps the best motives will be brought about by engaging in a course of psychotherapy, something that someone who had the right motives would never do.

If these points are not obvious then that may be because of a tendency to conflate the claim that certain motives are the right ones to have with the quite different claim that certain motives are the right ones *to try* to have: the right ones to try to inculcate or maintain in oneself. In estimating which motives it is best to have we must ignore the costs of inculcating or maintaining them: the evaluand is the *having of the motives*, not the *act* of getting to have them or keep them. However in estimating which motives it is best to try to have or keep we factor in such transition

costs: the evaluand is the *act* of getting to have or keep the motives, not just the *having of them*. The distinction is on a par with that to which Dr Johnson drew attention when he remarked that, yes, the Giant's Causeway was worth seeing – but that it was hardly worth going to see.

The only defence of the suggestion left, as far as we can see, is once again that the acts which are deemed to be right, on such a definition of rightness, accord more closely with our pre-theoretical intuitions about which acts are right and which are wrong than do the acts which are deemed to be right if we appeal directly to the consequentialist principle. But, once again, it seems to us that this requires that we forget about a whole range of possible cases in which people have the motives that it is right for them to have.

Consider, once again, the possible world in which there is a mad scientist who will make millions of people miserable if certain individuals don't have malignant motives, but who couldn't care less which acts these individuals perform. The malignant motives are, let's assume, the right motives for them to have in this possible world. Suppose that in this case, however, these individuals do not have these motives. It certainly doesn't accord with our pre-theoretical intuitions to suppose that the right way for them to act in this world is as if they had these malignant motives. Acting as if they had malignant motives will simply add to the disvalue that will be brought into the world by the mad scientist. It is much more plausible to suppose that they act rightly if they maximise value. The right thing for them to do is thus surely to act so as to acquire the malignant motives as quickly as possible, rather than acting as if they already had them.

The upshot, in our view, is thus once again that consequentialists who find themselves reluctant to believe that right acts are those which maximise value should reassess their entire commitment to a consequentialist way of thinking, not simply attempt to redefine rightness in the manner recommended by the second version of local motive consequentialism. For the idea that the right act for someone to perform is that act, of the acts available to them, which they would have performed if they had the right motives, unlike the idea that the right act is that which maximises value, is one which can be given no clear consequentialist rationale.

SECTION 5: LOCAL MOTIVE CONSEQUENTIALISM: THIRD VERSION

Consider finally the version of local motive consequentialism which holds that the right acts are those which would have been caused by the motives that it would be best for someone to try to inculcate. The main objection to this version, as in the last case, is that it is difficult to see what the consequentialist case could be for so acting.

Once again, note that there need not be the same non-act mediated benefits attached to actions that would have been caused by the motives it is best to try to inculcate as attach to actions that actually spring from such motives. The thought

must therefore be, as with the second version of local motive consequentialism, that acting as if you had the motives that it would be best for you to try to inculcate will generally produce or sustain those motives; that therein lies the local motive consequentialist case for so acting.

But this thought is once again confused. There are three relevant points. The first is that acting as if you had these motives might not in itself help you to inculcate them. The second is that if the right act is the act that the motives it is best to try to inculcate would cause, then it is not clear why it is right for anyone to try to inculcate them; the presence of the motives it is best to try to inculcate would not cause anyone, unnecessarily, to try to inculcate them. And the third is that the motives that you would inculcate in yourself by so acting might be manifestly inferior to others that you could have inculcated instead.

Nor is it plausible to suppose that this theory delivers a more intuitive conception of right action either. Consider the possible world in which there is a mad scientist who will make millions of people miserable if certain individuals don't have malignant motives in the future, but who couldn't care less which acts they perform now. Let it be agreed that it would therefore be best for them to try to have malignant motives. Yet it is surely quite implausible to suppose that these people act rightly, here and now, if they act as if they had these motives already. Once again, it is far more intuitive to suppose that they act rightly, here and now, if they act so as to maximise value; far more intuitive to suppose that it is right for them here and now to try to acquire the malignant motives, rather than acting as if they already had them.

It is therefore once again difficult to see what the consequentialist justification is supposed to be for performing the act which would have been caused by the motives that it would be best for someone to try to inculcate. The upshot is, again, that those consequentialists who find themselves reluctant to believe that right acts are those which maximise value should reassess their entire commitment to a consequentialist way of thinking rather than simply attempting to redefine rightness in the manner recommended by the third version of local motive consequentialism.

SECTION 6: GENERALISING THE ARGUMENT

We are now in a position to generalise our argument. Local rule consequentialism parallels fairly closely the sort of doctrine we have been describing as local motive consequentialism. Rules may refer to personal policies or to social practices and so local rule consequentialism may direct us to a local consequentialism at the individual or the collective level. But in either case the doctrine may assume any of the three forms described, and in either case it is subject to the same sorts of problems raised with local motive consequentialism.

Clearly, however rules are understood, local rule consequentialism may hold that

the right act is that which is caused by having internalised the rules it is best to have internalised; or that the right act is that which would be caused by having internalised the rules it is best for someone to have internalised; or that the right act is that which would be caused by having internalised the rules it is best for someone to try to internalise; or something of the kind. And clearly, so we maintain, familiar problems are going to arise for each of these versions of local rule consequentialism.

The first version has the absurd consequence that those who have not internalised the rules which it is best for them to have internalised cannot act rightly. But it is difficult to see what the consequentialist rationale might be for this claim when, among the acts available to them, some produce better consequences than others. The theory is also vulnerable on another score as well. For no consequentialist rationale can be given for the claim that those who have internalised the right rules always act rightly when their acts are caused by their internalisation of those rules. So long as there really is an alternative act available to an agent which produces more value, it is hard to see what there could be for a consequentialist to regret in the performance of such an act.

The idea that the first version of local rule consequentialism delivers a more intuitive conception of right action than that delivered by a direct evaluation acts in terms of the consequentialist principle is also refuted by a variation on the mad-scientist examples. Imagine a possible world in which there is a mad scientist who will make millions suffer if people don't internalise callous rules of conduct, but who couldn't care less which acts the people perform. Suppose that people all internalise these callous rules. Is it intuitive to suppose that these people act rightly when they act on their internalisation of these callous rules? Not at all. It is far more intuitive to suppose that they act rightly if they maximise value.

The second and third versions give rise to the same difficulty that plagued the second and third versions of local motive consequentialism. The question, again, is what the consequentialist case is supposed to be for redefining rightness in the ways proposed. Acting as if the rules that it is best to internalise or to try to internalise had been internalised will not have the same non-act mediated benefits associated with actually having internalised these rules. And acting in the required fashion cannot be justified on the grounds that it promises to bring about the internalisation of those rules either.

Here, as in the corresponding case with the second and third versions of local motive consequentialism there are three relevant points. The first is that acting as if you had internalised certain rules might not in itself help you to internalise the rules. The second is that if the right act is the act that having internalised certain rules would cause, then it is not clear why it is right for anyone to try to internalise them; the internalisation of those rules would not cause anyone, unnecessarily, to try to

internalise them. And the third is that the rules that you would internalise by so acting might be manifestly inferior to others that you could have internalised by acting in some alternative way instead.

The second and third versions of local rule consequentialism cannot plausibly be said to deliver a more intuitive conception of right action either. Variations on the mad-scientist examples once again suggest that, in certain circumstances, it is far more intuitive to suppose that agents act rightly when they maximise value. Consider the possible world in which there is a mad scientist who will make millions miserable if people haven't internalised callous rules of conduct here and now, and another possible world in which he will do so if they haven't internalised such rules in the future, but who couldn't care less how these people act here and now. Imagine further that people have not internalised the required callous rules of conduct here and now. It is surely quite counter-intuitive to suppose that they act rightly if they act as if they had internalised the callous rules here and now. For that will simply add to the misery that the mad scientist will produce without producing any clear benefit. It is far more intuitive to suppose that they act rightly if they act so as to maximise value; far more intuitive to suppose that the right thing for them to do is to internalise the callous rules, not to act as if they had already internalised them.

The situation is thus just as we said it would be at the outset. The very same pattern of objection that applies to the various versions of local motive consequentialism applies, as well, to the various versions of local rule consequentialism. We think that it is therefore reasonable to suppose that a similar pattern of objection would apply to all versions of local consequentialism. We readily admit that this requires an inductive leap, but the inductive leap required is, we think, hardly rationally resistible.

SECTION 7: CONCLUSION

Most philosophers now recognise that it is a mistake to think that right motives, right rules, and so on, are those that would promote right acts; they recognise that local act consequentialism, as we have called it, is unattractive. If the points we have made are sound, however, then they should also admit that it is equally unattractive to warp consequentialism in the manner illustrated by local motive or rule consequentialism. For going local in this latter way – defining right acts by reference to an independent account of right motives, right rules, or whatever – looks like it is always going to involve problems of the kind illustrated. There are many categories of evaluand: acts, sets of motives, sets of rules, as well as sets of such sets, and so on. If consequentialism is a sound strategy of evaluation in any one case, then it is hard to see why it should not represent a sound strategy in every case. Better, therefore, for consequentialists to go global; better for them to go the full monty.¹⁴

NOTES

1. The intuitive notion of rightness, as applied to acts, is thus just the idea of the act that it is right for someone to perform – as opposed, say, to the act that it is right for someone to forget. In the latter case the evaluand is forgettings, not acts.
2. Derek Parfit, *Reasons and Persons* (Oxford: Clarendon Press, 1984), pp. 24–5.
3. J. J. C. Smart appears to advocate a form of global consequentialism as well when he argues in favour of what he calls 'extreme utilitarianism'. See note 6 below.
4. R. M. Adams, 'Motive Consequentialism', *Journal of Philosophy*, 73 (1976), pp. 467–81.
5. 'An Outline of a System of Utilitarian Ethics', in J. J. C. Smart, and Bernard Williams, *Utilitarianism: For and Against* (Cambridge: Cambridge University Press, 1973), p. 9.
6. Immediately after the passage just quoted Smart of course goes on to argue against rule utilitarianism. Interestingly, in the earlier paper in which Smart argues in favour of a form of global consequentialism he also appears to argue against local consequentialism more generally, though he argues against it under the name 'restricted utilitarianism'. See his 1956 paper 'Extreme and Restricted Utilitarianism', reprinted in Philippa Foot (ed.), *Theories of Ethics* (Oxford: Oxford University Press, 1967).
7. Brad Hooker, *Ideal Code, Real World: A Rule-consequentialist Theory of Morality* (Oxford: Clarendon Press, 2000).
8. Adams, 'Motive Consequentialism', p. 479.
9. Simon Blackburn, *Oxford Dictionary of Philosophy* (Oxford: Oxford University Press, 1994), p. 388.
10. Adams, 'Motive Utilitarianism'; Peter Railton, 'Alienation, Consequentialism, and the Demands of Morality', *Philosophy and Public Affairs*, 13 (1984), pp. 134–71; Parfit, *Reasons and Persons*.
11. E.g., Adams, 'Motive Utilitarianism'.
12. Note that the restrictive form of consequentialism advocated by Pettit and Brennan is consistent with global consequentialism ('Restrictive Consequentialism', *Australasian Journal of Philosophy*, 64 (1986), pp. 438–55). According to restrictive consequentialism, even though the right act is that which maximises value, the right decision-making procedure – the procedure which maximises value – may not be to calculate case by case in an attempt to identify the act that maximises value. Rather, as Pettit and Brennan argue, it might instead be better to restrict one's calculation. 'Restrictive' in the sense of Pettit and Brennan is thus not synonymous with 'restricted' in the sense defined by Smart's 'Extreme and Restricted Utilitarianism'.
13. Fred Feldman, 'World Utilitarianism', in Keith Lehrer (ed.), *Analysis and Metaphysics* (Dordrecht: Reidel, 1975), pp. 255–71.
14. We would like to thank Lloyd Humberstone and Vera Koffman for their helpful remarks on an earlier version of this chapter.