



Good and Bad Actions

Author(s): Alastair Norcross

Source: *The Philosophical Review*, Vol. 106, No. 1, (Jan., 1997), pp. 1-34

Published by: Duke University Press on behalf of Philosophical Review

Stable URL: <http://www.jstor.org/stable/2998340>

Accessed: 15/04/2008 17:24

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=duke>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit organization founded in 1995 to build trusted digital archives for scholarship. We enable the scholarly community to preserve their work and the materials they rely upon, and to build a common research platform that promotes the discovery and use of these resources. For more information about JSTOR, please contact support@jstor.org.

Good and Bad Actions

Alastair Norcross

It is usually assumed to be possible, and sometimes even desirable, for consequentialists to make judgments about both the rightness and the goodness of actions. Whether a particular action is right or wrong is one question addressed by a consequentialist theory such as utilitarianism. Whether the action is good or bad, and how good or bad it is, are two others. I will argue in this paper that consequentialism cannot provide a satisfactory account of the goodness of actions, on the most natural approach to the question. I will also argue that, strictly speaking, a consequentialist cannot judge one action to be better or worse than another action performed at a different time or by a different person. Even if such theories are thought to be primarily concerned with rightness, this would be surprising; but in the light of recent work challenging the place of rightness in consequentialism,¹ it seems particularly disturbing. If actions are neither right (or wrong) nor good (or bad), what moral judgments do apply to them? Doesn't the rejection of both rightness and goodness, as applied to actions, leave consequentialism unacceptably impoverished? On the contrary, I will argue that consequentialism is actually strengthened by the realization that actions can only be judged as better or worse than possible alternatives.

I am grateful to many people for help with this paper, some of whom I cannot name. Earlier versions of the paper were presented to audiences at the North Texas Philosophical Association, the American Philosophical Association Pacific Division, the University of Houston, the University of Arkansas, and Texas Tech University. I received many helpful comments and suggestions from members of those audiences, some of whose names I have, unfortunately, forgotten. I have also received many helpful comments from my colleagues at SMU, whose names I couldn't forget, even if I tried (which I have no desire to do). I have received extensive and very helpful comments on earlier versions of this paper from several anonymous referees for the *Philosophical Review*. Those I would like, and am able, to thank by name are Michael Almeida, Eric Barnes, Philip Clark, Douglas Ehring, David Hausman, Mark Heller, Christopher Hitchcock, Frances Howard-Snyder, Jean Kazez, James Lamb, David Phillips, Tom Senor, Michael Slote, Steven Sverdlik, Peter Vallentyne, and Mark Webb.

¹See Slote 1985a, chap. 5, and Howard-Snyder and Norcross 1993.

1. Goodness and Rightness

Consequentialism has traditionally been viewed as a theory of right action. Consequentialists have employed theories of value, theories that tell us what things are good and bad, to provide inputs for functions whose outputs tell us what actions are right and wrong. The theory of the good is usually taken to be a theory of the goodness of states of affairs. The most common consequentialist function from the good to the right embodies the maximizing requirement:

R An act is right iff there are no available alternatives that produce a greater balance of goodness over badness.

That is, the right action is simply the best action.² So, what should consequentialism say about *good* actions? What, if any, is—or, rather, should be—the connection between the consequentialist accounts of right actions and good actions?

There are three different approaches to these questions with *prima facie* appeal, which can be roughly characterized as follows: (i) ‘right’ and ‘good’, as applied to actions, are interchangeable, except for the fact that ‘good’ admits comparative and superlative forms; (ii) the goodness of an action is a function of the goodness of the motive (or maybe even the whole character) from which it sprang; (iii) the goodness of an action is a function of the goodness of its consequences. The last of these seems to me the most natural approach for a consequentialist to adopt, and so the argument of this paper, apart from the brief remarks of the next two paragraphs, is directed towards (iii). I leave open the possibility that a consequentialist account of good actions along the lines of (i) or (ii) could be provided, although, as I will now explain, even such accounts will be affected by my arguments regarding (iii).

Option (i) has a certain intuitive appeal, rooted in the ordinary usage of the terms ‘right’ and ‘good’ beyond a merely consequentialist framework, though our intuitions (at least mine and those of others I have consulted) are by no means univocal on this point. The upshot of adopting (i), for a consequentialist, is that all and

²For convenience, I will ignore the possibility of ties.

only right actions are good actions, and that some good actions are better than others. This is particularly counterintuitive on the most popular consequentialist account of rightness, *R*. I will have more to say about a maximizing account of good actions in section 2. Furthermore, how should a maximizing consequentialist understand the claim that one good/right action is better than another? The obvious answer seems to be that the former action leads to a greater balance of good over bad³ than does the latter. However, I will argue, in sections 2–4, that a consequentialist cannot give a satisfactory account of the notion of an action's leading to a balance of good over bad. I will also argue, in section 5, that it follows from this that a consequentialist cannot say, strictly speaking, that one actual action is either better or worse than another. These difficulties also affect the attempt to tie (i) to a satisficing account of rightness/goodness, on the most intuitive reading of such an account.⁴

Option (ii) has been suggested by both consequentialists and nonconsequentialists as an account of good actions. On the non-consequentialist side we have W. D. Ross:

Now when we ask what is the general nature of morally good actions, it seems quite clear that it is in virtue of the motives that they proceed from that actions are morally good. (1973, 156)⁵

A notable consequentialist who takes a similar type of approach is J. J. C. Smart:

We can also use 'good' and 'bad' as terms of commendation or dis-

³Or smaller balance of bad over good.

⁴Such an account would specify that in order to be right/good, an action would have to produce a certain balance of good over bad. Other satisficing accounts might demand that an action be no further than a certain distance from the best option, in order to be right/good. This might be a fixed distance, or it might vary with the context of the choice, or even of the evaluation itself. This type of approach, except for the variation that has the rightness/goodness of actions vary with the context of evaluation (see section 6), is not affected by the argument of this paper. However, see Howard-Snyder and Norcross 1993 for arguments in favor of scalar consequentialism over satisficing consequentialism.

⁵Ross sees his account as being in the tradition of that most famous of nonconsequentialists, Immanuel Kant.

commendation of actions themselves. In this case to commend or discommend an action is to commend or discommend the motive from which it sprang. (1973, 49)

The first problem that springs to mind with this approach is the problem of individuating motives. For example, take the case of the mother of a vicious killer who lies to the police to keep her son out of jail, despite knowing exactly what he did. Consider the following two, equally plausible, accounts of her motive: (a) she lied because she wanted to protect her child from harm; (b) she lied because she wanted to protect her child from the consequences of his own terrible wrongdoing. Whether her lie is good or bad could depend, according to Smart's approach, on which account of her motive we accept. The problem here is not just an epistemic one. It's not as if we just don't know whether her motive was really to protect her child from harm or to protect him from the consequences of his own terrible wrongdoing, but if we did could evaluate the motive. To ask which of the alternatives was really her motive might be like asking whether it was really the President of the United States or Bill Clinton who gave the 1996 State of the Union Address. However, let's assume that we have a satisfactory method for individuating motives. The next question for the proponent of this approach is whether the motive is evaluated with reference to the specific agent or a broader class of agents. A motive that in most people leads to good results could lead to terrible results in a few. Finally, what evaluation of the motive is required in order for the action to be good? Must the motive (in this agent or some wider class of agents) lead to a mere balance of good over bad, a particular positive balance of good over bad, the greatest possible balance of good over bad?⁶ Only the maxi-

⁶Smart seems to make the mistake of assessing the motive in terms of the right actions to which it leads. As an explanation of why we should approve of the desire to save life, even though it sometimes leads to doing the wrong thing, he says, "[I]n general, though not in this case, the desire to save life leads to acting rightly" (1973, 49). But clearly a consequentialist would compare motives in terms of the goodness of their consequences, not the rightness. Consider two possible competing motives, A and B, which are relevant to three different choices, each between two different actions. For the first two choices, A would produce the better action and B the worse action. For the last choice, B would produce the better action and A the worse action. So, in this simplified context, A would lead to more right actions (on a maximizing theory such as Smart's) than would

mizing alternative will be unaffected by my treatment of option (iii), since my arguments against the notion of an action's leading to a balance of good over bad can be adapted to apply to the notion of a motive's leading to a balance of good over bad. I leave open the possibility, therefore, however distant, that some account of good actions could be produced that ties the goodness of an action to some consequential feature of the motive from which it sprang.

Let's return to (iii): the goodness of an action is a function of the goodness of its consequences. How might a consequentialist employ her value theory to give a theory of good actions? It doesn't follow from the fact that we have a theory of the goodness or badness of states of affairs that we have a theory of the goodness or badness of actions, but there might seem to be an easy method for constructing the latter out of the former:

G An act is good iff it produces more goodness than badness;
an act is bad iff it produces more badness than goodness.

So far as I know, no consequentialist has advocated *G* in print,⁷ though the central concept, that of an action's producing a balance of goodness over badness, is either explicit or implicit in much consequentialist literature. Indeed, if we consider what maximizing consequentialism tells us about assessing rightness, the procedure seems to be as follows: first, determine how much goodness and badness each possible act will produce; next, rank them according to the sum, positive or negative, of goodness over badness; finally, declare the act with the greatest sum to be the right act. Thus, Sidgwick:

By Utilitarianism is here meant the ethical theory, that the conduct which, under any given circumstances, is objectively right, is that which

B. Should a consequentialist prefer A to B (assuming that the same motive has to be operative in each choice)? That depends on the details of the three choices. If the better, and therefore right, action in the first two choices is only slightly better than the alternative, but the better action in the last choice is much better than the alternative, the consequentialist would clearly prefer motive B over A.

⁷Though Michael Slote suggested it in comments on an early version of Howard-Snyder and Norcross 1993.

will produce the greatest amount of happiness on the whole; . . . by Greatest Happiness is meant the greatest possible surplus of pleasure over pain, the pain being conceived as balanced against an equal amount of pleasure. (1981, 411, 413)

If each act produces either goodness or badness (or both), and it is possible for amounts of goodness and badness to cancel each other out, we ought to be able to classify acts as either good, bad, or neutral, according to whether the sum of goodness over badness is positive, negative, or neither. This possibility is suggested by Bentham's claim that

[a]n action . . . may be said to be conformable to the principle of utility . . . when the tendency it has to augment the happiness of the community is greater than any it has to diminish it. (1789, chap. 1, para. 6)

Bentham is usually understood to be talking about rightness.⁸ This would seem to be a nonmaximizing account of rightness that locates the threshold between right and wrong actions at the same point as the threshold between good and bad actions, according to *G*. A combination of *G* with Bentham's account of rightness could give us a version of (i), discussed above: all and only right actions are also good actions, but some are better than others.

It also seems possible simply to equate rightness with goodness (and wrongness with badness), according to *G*. In fact, Mill's statement of the principle of utility can be interpreted as suggesting such a scalar interpretation of rightness (and wrongness):

The creed which accepts as the foundation of morals, Utility, or the Greatest Happiness Principle, holds that actions are right in proportion as they tend to promote happiness, wrong as they tend to produce the reverse of happiness. (1861, chap. 2)

Another possibility is to incorporate both the maximizing requirement and the distinction between good and bad actions in an account of rightness. Slote suggests a modification of utilitarianism to

⁸See, for example, Quinton 1989, 1–3.

the moral theory that results if one demands of a right action both that it produce consequences no less good than those producible by any alternative act available to a given agent and that those consequences be, on balance, good. (1985b, 162)

The distinction between good and bad actions, judged from a consequentialist point of view, might even figure in a nonconsequentialist theory of right action. Shelly Kagan describes a theory that incorporates a “zero threshold constraint against doing harm” (1989, 191ff.). Here’s some of what he has to say about this constraint:

the zero threshold constraint against doing harm forbids doing harm in those cases where this will lead to an overall loss in objective good. It does not, however, provide any barrier to doing harm in those cases where this will result in an overall gain. . . . [S]o long as the harm brings about *better* consequences overall, it need not bring about the *best*. For so long as the suboptimal act of harm-doing will, on balance, bring about more good than harm, it will not be ruled out by the zero threshold constraint. (1989, 191–92)

None of these philosophers is here explicitly advocating a theory of good actions, but they all seem to be using something like the concept involved in *G*, that is, the concept of an action’s producing a balance of goodness over badness. It would appear that a consequentialist can classify as good any action that will *on balance, bring about more good than harm*, or whose *consequences are, on balance, good*, or whose *tendency to augment the good of the community is greater than any it has to diminish it*. Indeed, these phrases appear to capture a single common and commonly accepted consequentialist concept. But that appearance is deceptive. I will argue that there are, in fact, several concepts that they might capture, but that since none of these yields a plausible version of *G*, there is no satisfactory way for a consequentialist to use *G*, or anything like it, to judge actions as simply good or bad, as opposed to better or worse than specific alternatives.

2. Goodness and Comparisons

How might we explain what it is to augment the good of the community, or for the consequences of an action to be, on balance,

good? For the sake of simplicity, I will assume happiness and unhappiness to be the only things of intrinsic value and disvalue. Consider an agent, called Agent, whose action affects only herself and one other person, Patient. Agent is faced with a range of options that do not affect her own happiness, but that have dramatically different effects on Patient's happiness. This case seems simple enough. The good actions are those that make Patient happy, the bad are those that make him unhappy.

But this won't yet do. It seems to assume that Patient was neither happy nor unhappy to begin with. Let's modify the account slightly. The good actions are those that make Patient *happier*, the bad are those that make him *unhappier*. Happier than what? One obvious answer is happier than he was before the action. If Agent does something that increases (or augments) Patient's happiness, she has done a good thing. To generalize, we simply compare the welfare of all those affected by a particular action before and after the action. If the overall level of welfare is higher after than before, the action is good. If it is lower, the action is bad. If it is the same, the action is neutral.

But this still won't do. Consider again a restricted case involving only Agent and Patient. Call this case *Doctor*. Patient is terminally ill. His condition is declining, and his suffering is increasing. Agent cannot delay Patient's death. The only thing she can do is to slow the rate of increase of Patient's suffering by administering various drugs. The best available drugs completely remove the pain that Patient would have suffered as a result of his illness. However, they also produce, as a side effect, a level of suffering that is dramatically lower than he would have experienced without them but significantly higher than he is now experiencing.⁹ So the result of ad-

⁹The rate of increase of pain is essential to the example. It is important that Patient suffer more after the treatment than before, because the view I am arguing against involves a simple comparison of Patient's welfare before and after the action. If the treatment left Patient in less pain after the action than before, it would count as a good action, both intuitively and according to the account in question. Why not simply have an example involving a regular painkiller that removes some, but not all, of the pain that Patient would have suffered? It would avoid possible complications if Patient's later states are uncontroversially caused by Agent's action. According to some intuitively appealing accounts of causation and mental-state identity, ordinary painkillers do cause the later painful states. However, there are theories of causation and mental-state identity according to

ministering the drugs is that Patient's suffering continues to increase, but at a slower rate than he would have experienced without them. The very best thing Agent can do has the consequence that Patient's suffering increases. That is, after Agent's action, Patient is suffering N amount of suffering as a direct result of the action, and N is more than Patient was suffering before the action. Has Agent done a *bad* thing if she slows the rate of increase of Patient's suffering as much as she can? This hardly seems plausible. It is consistent with the schematic description of this case to imagine that Agent has done a very good thing indeed.

Clearly, we can't simply compare states of affairs before and after a particular action. Agent *has* made Patient happier: not happier than he was, but happier than he would have been. We compare states of affairs, not across times, but across worlds. Agent has done a good thing, because she has made Patient happier than he would have been had she done something else. Even though Patient is now suffering more than he was, he would have been suffering even more if Agent had done anything else instead.

As I said, an evaluation of Agent's action involves a comparison of different worlds. But which world (or worlds) do we compare with the world containing Agent's action?¹⁰ With what do we com-

which this is not so, and I don't want my example to depend on the truth of any particular controversial metaphysical view.

¹⁰For the sake of simplicity, I pretend throughout this paper that the world is deterministic. Thus, I talk of *the world* in which the action is performed. Perhaps, though, an acceptance of indeterminism will provide a method of assessing the goodness or badness of actions. Consider the following sketch of an account, suggested both by an anonymous referee and by Mark Brown: A possible action determines a cone of worlds: all possible histories of the universe coinciding with the actual world up to the point of action and in which the action gets done. The value of the action, in terms of goodness or badness, is that of its cone. The value of the cone is determined by the value of the post-act part of the worlds in it, probably by integrating their value weighted by their probability of being the actual world. Unfortunately, the example of *Doctor* can be easily modified to show this kind of approach to be inadequate. Imagine that Agent and Patient are the last two living sentient beings, and that Agent is suffering from a similar condition to Patient's, but less advanced. After they are both dead, there will be no more morally relevant beings, ever. Further imagine that Agent is unable to kill either herself or Patient. The best she can do for Patient is to administer the drugs that result in a slowed rate of increase of Patient's suffering. It is highly plausible that the action cone for Agent's act of administering the drug to Patient is bad. The probability of a miraculous recovery for either Agent or Patient is negligible. All, or almost all, of the post-act portions of the worlds in the action cone are bad. So,

pare Patient's suffering? There are several other ways for Agent to behave. Which of these alternatives provides the relevant comparison? In this case it doesn't seem to matter, because Agent has done the best thing possible. None of her other available options would have resulted in less suffering for Patient than did her actual behavior. But we don't want to demand of a good action that there be no better alternatives. This would be a maximizing account of good actions that would equate goodness with the maximizing notion of rightness. To see why this is unacceptable, consider the following case, *Self-sacrifice*. Agent is able, at the cost of some considerable effort and pain to herself, to make Patient moderately happy. She does so. This is nearly the best thing that she could do, but not quite. One alternative course of action would have made Patient considerably happier, while all her other alternatives would have resulted in far less happiness for him, and some would even have led to unhappiness. The action that would have resulted in more happiness for Patient, however, would have involved a fair bit more effort and pain for Agent. The extra effort and pain for Agent would have been slightly outweighed by the extra happiness for Patient, but only slightly. So there is something that Agent could have done that would have had even better, albeit marginally better, consequences than what she did. But do we really want to say that Agent didn't do a good thing in sacrificing her own comfort for the sake of Patient's happiness? She didn't do the *best* thing,

any plausible method of integrating the value of the worlds in the cone will yield the result that the cone is, overall, bad. Of course, all the action cones of the alternative possible actions will be even worse. But it is highly implausible to suggest that Agent's action is bad but not as bad as the alternatives. If we have an account of good and bad actions, it should judge her action to be good. More generally, the problem with the action cone approach is that it judges many intuitively good actions to be bad and many intuitively bad actions to be good. Consider a situation in which it's overwhelmingly likely that no matter what I do, goodness will outweigh badness throughout the future of the world. The worst thing that I can do is to torture and kill five people. Even if I do that, however, there will probably be such an abundance of goodness throughout the rest of the world, maybe in terms of happiness, that the post-act portions of the worlds in my action cone will be, almost exclusively, good. The integrated value of my action cone is good. But my action is clearly not good. My action is judged as good, on this approach, because the goodness of states of affairs that are unaffected by it outweighs the badness I bring about. This suggests a modification of this approach that counts only those states of affairs that are affected by an action. I discuss such an approach in section 4.

but what she did seems to be pretty good. This is not to say, of course, that there are no situations in which only the best action is plausibly regarded as good, whether by consequentialists or others. There may even be situations in which great self-sacrifice is required in order to do good (the example of *Lifeboat*, discussed below, is possibly such a case). Such cases will typically involve the prevention of some great harm to another. In *Self-sacrifice*, however, Agent is already providing Patient with a considerable benefit, at no small cost to herself. The only motivation I can see for insisting that her action is not good, on the grounds that she can do even better, is the determination to equate the notions of the right and the good as applied to actions. If we are to give a consequentialist account of good actions, we should accommodate the intuition that at least some suboptimal acts are nonetheless good.

Self-sacrifice demonstrates that optimization is not an appropriate standard of goodness, but it also suggests a different approach. The reason why Agent doesn't have to optimize in order to do good, it might be claimed, is that optimization involves, in this case, a greater sacrifice of her own interests than is required for mere goodness. (This leaves open the possibility that optimization is nonetheless required for rightness.) At this point we might be tempted to adapt Samuel Scheffler's agent-centered prerogative (1982) to apply instead to goodness.¹¹ Consider the following account:

GAC An act is good iff either (i) it is optimal,¹² or (ii) producing better consequences would require showing less than a certain proportionate bias toward consequences for the agent.

According to *GAC*, Agent's action in *Self-sacrifice* could still be good, so long as the better action would have required showing less than the relevant degree of bias towards herself. It seems plausible to assume that whatever the relevant degree of bias is, it will be greater than Agent would show in performing the best action. This is because pretty much *any* bias towards herself would result in Agent preferring the second-best over the best action. Recall that the

¹¹I owe this suggestion to a referee for the *Philosophical Review*.

¹²Without (i), *GAC* would give the strange result that many suboptimal acts were good while their optimal alternatives were not.

extra burden that the best action would have imposed on Agent would only just have been outweighed by the extra benefit for Patient.

Despite the success of *GAC* in coping with *Self-sacrifice*, I don't think consequentialists should embrace it as an account of good actions. There are two reasons for this. First, *GAC* is an agent-relative account of good actions. The classical utilitarians all endorsed a non-agent-relative standard for assessing actions. Most famously, Bentham required "everybody to count for one, and nobody for more than one," and Mill said of the utilitarian agent, "As between his own happiness and that of others, utilitarianism requires him to be as strictly impartial as a disinterested and benevolent spectator" (1861, chap. 2). These claims were made in connection with assessing the rightness of actions, but they embody a central feature of consequentialist ethical theories. Scheffler's agent-centered prerogative is seen as a departure from consequentialism, not simply because it rejects maximization, but because the rejection of maximization is achieved by allowing agents a degree of partiality toward themselves. Those consequentialists for whom the disinterested benevolent spectator provides the appropriate model of moral assessment of actions will be loath to abandon that model when it comes to assessing the goodness of actions. I don't wish to claim that it would be inappropriate to call a view incorporating agent-relative standards "consequentialist." I suspect that non-agent-relativity in all action judgments is part of what distinguishes consequentialist theories from other ethical theories, but that is the topic of another paper.¹³ At the very least, it is worth noting that even if *GAC* were otherwise acceptable as an account of good actions, it would be unappealing to those consequentialists who embrace non-agent-relativity in all action judgments. But *GAC* is not otherwise acceptable.

Optimization is unacceptable as the standard of goodness, because it excludes too much, such as Agent's action in *Self-sacrifice*. *GAC* expands the realm of good actions to include this action and others like it. However, *GAC* both excludes too much and includes too much, as the following examples demonstrate. In *Self-sacrifice 2*, Agent is able, at the cost of some considerable effort and pain to herself, to make Patient moderately happy. She does so. This is

¹³For a useful discussion of this issue, see Howard-Snyder 1994.

nearly the best thing that she could do, but not quite. Her best option involves shifting the burden of making Patient happy from herself onto a third person, Other. The pain and sacrifice for Other in the best option would have been fractionally less than Agent bore in the second-best option. The happiness for Patient would have been identical. Agent's action is not optimal, nor would producing better consequences require her to show less than the requisite amount of bias towards herself, since the best action—shifting the burden to Other—involves far more bias towards herself than does the second-best action. According to GAC, therefore, Agent's action of bearing the burden of making Patient happy rather than imposing a fractionally smaller burden on someone else is not good. But this is highly implausible. How could a supporter of GAC defend this result? Perhaps she could argue as follows: Optimization is the default standard of goodness. However, agents are permitted a certain bias towards themselves. Thus, some departures from optimization can still be good, if better actions would have required excessive self-sacrifice. This is what allows Agent's action in *Self-sacrifice* to count as good. In *Self-sacrifice 2*, Agent can do better at *less* cost to herself, so her decision to bear the burden of making Patient happy is just pointless masochism. I have two replies to this. First, why should optimization be the default standard of goodness? Aren't there many nonoptimal actions that are intuitively good, even when better actions wouldn't involve significant self-sacrifice? Isn't my giving \$50 to a worthy charity a good action, even though I could have given \$51 without significant self-sacrifice? Second, Agent's action in *Self-sacrifice 2* is not pointless masochism. The point is to spare Other the burden of helping Patient. This becomes clearer when we fill in the details of the case. Suppose my young child is sick and miserable and needs comfort in the middle of the night. My wife could provide the same amount of comfort as I at fractionally less cost to herself (she has a slightly less burdensome day ahead). Nonetheless, I drag myself out of bed and let her get a good night's rest.

GAC also classifies as good some actions that clearly aren't, such as in the following example, *Lifeboat*. Agent and Patient are adrift in a lifeboat, with only enough food to sustain one of them until help arrives. If they attempt to share the food, they will both die. Agent is a second-hand car salesman, who specializes in selling lemons. Patient is a dedicated physician, who runs a free clinic for

poor children in the inner city. While Patient is sleeping, Agent tips her over the side of the boat, thus ensuring his own survival. This is not the best action, but to do better, Agent would have had to have sacrificed his own life. Such a sacrifice would clearly have involved showing less than the permitted bias towards himself. According to *GAC*, then, Agent's action of tipping Patient over the side of the boat is good. But this won't do. However excusable we may deem Agent's action to be, it is by no stretch of the imagination good.

3. Goodness and Counterfactuals

When we think of someone doing a good or a bad thing, an underlying concept, I suggest, is that of making a difference to the world. It is natural to think of a good action as one that makes the world better than it would have been if the action hadn't been performed. This suggests the following interpretation of *G*:

GC An act *A* is good iff the world would have been worse if *A* hadn't been performed; *A* is bad iff the world would have been better if *A* hadn't been performed.

This gives the right result in *Doctor*. If Agent hadn't done what she did, the world would have been worse. Patient would have been suffering even more. *GC* assesses the goodness of an action by comparing the world in which it occurs with a world in which it doesn't occur. But which world in which it doesn't occur is the relevant one? In *Doctor*, for example, there are many different ways that Agent could have failed to administer the pain-reducing drugs to Patient. There are many different things that she could have done instead, including doing nothing. In this case, we don't have to know just what Agent would have done instead, because we know that she did the best she could; so *anything* else would have been worse.

The intuitive reading of *GC* involves a comparison with the world in which the agent is inactive. When we ask what the world would have been like if the action hadn't been performed, we are considering a world in which the agent simply doesn't exercise her agency. So, what is it not to exercise one's agency? One obvious possibility is that it is to remain completely immobile. But this clear-

ly won't do. Consider the following case, *Button pusher*. Agent stumbles onto an experiment conducted by a twisted scientist named Scientist. He is seated at a desk with ten buttons, numbered '0' through '9', in front of him. He tells Agent that the buttons control the fates of ten people. If no button is pressed within the next thirty seconds, all ten will die. If the button marked '9' is pressed, only nine will die; if '8' is pressed, eight will die, and so on down to '0'. He was, he explains, about to sit and watch as all ten died. However, to honor her arrival, he turns control of the buttons over to Agent. She is free to press any button she wishes, or to press none at all. Agent pushes '9', killing nine people. If she had remained immobile, all ten would have died. According to *GC*, then, her action is good, since the world would have been worse, if she hadn't performed it. But her action is not good. It led to the deaths of nine people who needn't have died. She could have pressed '0' instead. Any satisfactory account of good actions has to judge this to be a bad action.

Perhaps we should compare the results of Agent's action with what would have happened if Agent hadn't even been on the scene. There seem to be two ways to interpret this suggestion: (i) we imagine a world identical to the actual world before *t*, in which the agent miraculously vanishes from the scene at *t*; (ii) we imagine a world as similar as possible to the actual world before *t*, in which the agent is nonmiraculously absent from the scene at *t*. That is, we imagine what would had to have been different before *t* in order for the agent to have been absent at *t*. Interpretation (i) runs afoul of *Button pusher*. If Agent had miraculously vanished instead of pushing button '9', all ten people would have died. But this consideration clearly doesn't incline us to judge that Agent did a good thing. Interpretation (ii) seems more promising. How do I know whether I have done a good thing? I ask myself whether I have made things better than they would have been had I not even been here in the first place. But this won't do, either. Once again, it gives the wrong results in *Button pusher*. If Agent hadn't even shown up in the first place Scientist would have let all ten die, but we don't on that count judge Agent's action to be good.

In *Button pusher*, Agent kills nine people, but ten would have died had she been inactive, either through immobility or absence from the scene. The problem is not just that inactivity gives unacceptable results in particular cases, but rather that the comparisons it invites

do not seem relevant to the goodness or badness of actions. If I do something that seems to be very bad, such as killing nine people with the press of a button, why should it matter that the consequences would have been even worse if I had been immobile or absent from the scene? Whether it is a good thing or a bad thing to kill nine people doesn't seem to depend on whether even more would have died if I had been inactive, unless, perhaps, my killing nine is the only alternative to more deaths. In *Button pusher*, however, Agent could easily have prevented any deaths. These counterfactuals, then, don't seem relevant to the goodness or badness of the actions, and they certainly don't capture what we mean when we ask what would have happened if the agent hadn't exercised her agency.

There are other ways to read the counterfactuals in *GC* that will give different accounts of goodness and badness. The most obvious alternative reading involves a judgment about which other possible world is closest to the world in which the action occurs.¹⁴ Instead of comparing the world in which the action occurs with a world in which the agent is either immobile or absent from the scene, we compare it with a world that is as much like it as possible, consistent with the action not occurring. Sometimes that will be the world in which the agent is immobile, but often it will be a world in which the agent does something else instead. Let's say I hit the bull's eye while playing a game of darts. Given that I'm not a particularly good darts player, the closest world in which I don't hit the bull's eye is probably one in which I just miss it and hit the collar around it instead. According to this interpretation of *GC*, we consider what the agent would have done instead if she hadn't performed the action in question.

How does this interpretation of *GC* handle *Self-sacrifice* and *Button pusher*? Consider *Self-sacrifice* first. Would the world have been worse if Agent hadn't done what she did? Since she didn't do the best possible thing, that depends on what she would have done instead. What Agent did in *Self-sacrifice* was very nearly the best. It involved a considerable amount of effort and pain for herself. It might seem plausible to assume, then, that if she hadn't done that, she would have done something worse. A deviation from her actual action that required less self-sacrifice would have been easier, and thus

¹⁴See, for example, Lewis 1973.

more likely, than one that required more. A world in which she sacrifices less would seem to be closer to the actual world than would one in which she sacrifices more. But these are guesses, based on a very sketchy description of the case. Let's add a couple of details to my previous description of *Self-sacrifice*. Agent is a committed consequentialist, highly predisposed toward self-sacrifice. There are no options available to her that are only slightly worse than what she did. In fact, the next best thing that she could do is much worse, involving suffering for Patient. The addition of these details makes it much less likely that Agent would have done something worse if she hadn't done what she did. In fact, it now seems overwhelmingly likely that she would have done even better if she hadn't done what she did. According to *GC*, then, her action was bad, since the world would have been better if it hadn't been performed. But her action wasn't bad. The fact that Agent's character made it extremely unlikely that she would have done worse than she did doesn't alter our intuitive judgment that her action was good.

Now for *Button pusher*. Once again, let me add some details to my previous description of the case. Agent is highly misanthropic. She delights in the misfortunes of others, especially their deaths. Her initial inclination is to refrain from pushing any buttons, so that all ten will die. She is dissatisfied, though, that this will involve, as she sees it, merely letting people die. She wants as many as possible to die, but she also wants to kill them. At the last second she changes her mind, and pushes '9'. If she hadn't pushed '9', she wouldn't have pushed any button. She didn't even consider the possibility of pushing a different button. The only question she considered was whether she should kill 9 or let 10 die. Clearly, the closest world in which Agent doesn't push '9' is one in which she doesn't push any button and all ten die. Once again, *GC* judges Agent's action to be good. But we are no more inclined to believe that her action is good than we were before we knew about her character defects. The fact that Agent's character made it highly probable that she would have done even worse than she did doesn't alter our intuitive judgment that her act of killing nine people was bad.

What is particularly disturbing for a consequentialist about this latest reading of *GC* is that it makes the character of the agent

relevant to the goodness of the action.¹⁵ The better the agent, the harder it is for her to do something good, and the worse she is, the easier it is. I don't deny that the character of agents can influence *some* judgments about actions or other events in a fairly natural way. For example, our judgments about whether it's a good thing that something happened are often influenced by our prior expectations. If we would have expected an outcome of an event or action to be worse, whether because of our knowledge of the characters of the agents involved or because of our prior experience of similar events, we may be pleased to discover that things aren't as bad as they might have been. Thus, we might claim that it's a good thing that only ten people were killed in the plane crash, or that the Republican Congress cut entitlements by only 80 percent. We don't mean by this that the budget cut was good, just that it wasn't as bad as we were expecting.

It seems that none of the interpretations of *GC* can provide the consequentialist with a satisfactory account of what it is for an action to be good. The intuition on which they were based is that a good action makes the world better. The difficulty lies in producing a general formula to identify the particular possible world (or worlds) than which the actual world is better as a result of a good action. Any unified theory requires a way of fixing the contrast point, but the contrast point varies from situation to situation. Part of the problem is that our intuitions about the goodness or badness of particular actions are often influenced by features of the context that it would be difficult to incorporate into a general account. I will have more to say about this in section 5, where I will consider a counterfactual account that isn't general.

4. In Search of a Noncomparative Account

I began this paper with some examples of philosophers who seemed to be making use of the concept of an action's consequences being, on balance, good. I have examined various interpretations of this concept and found that none provides a plausible

¹⁵The problem here is both that the proposal makes character relevant to the goodness of actions at all and that it does so in a particularly counterintuitive way. For a consequentialist, the former problem is more significant.

consequentialist account of what it is for an action to be good. Each has involved a comparison between the world that results from the action and a different, uniformly specified, possible world. It might be objected at this point that this is the wrong approach. If I claim that an action's consequences are, on balance, good, I am making a claim about the absolute value of certain states of affairs, not about their comparative value. If I am a hedonist, for example, I am not claiming that the world that resulted from the action contained a greater balance of pleasure over pain than any other particular world, I am claiming that the consequences of the action contained a balance of pleasure over pain. The fact that none of the comparative accounts of an action's goodness proved satisfactory is entirely to be expected. We should be looking for a noncomparative account.

If we take this objection seriously, we must ask what are the consequences of an action. Recall the scenario in *Doctor*. Agent administers a drug to Patient, who endures a great deal of suffering before eventually dying. Is Patient's suffering a consequence of Agent's action? It is hard to see how it couldn't be. The drug that Agent administered produced the suffering. Patient experiences a good deal of pain and no pleasure as a result of Agent's action. Assume that Agent experiences neither pleasure nor pain. Agent and Patient are the only people involved. It would seem that the consequences of Agent's action contain a balance of pain over pleasure. But Agent is doing the best she can; she is slowing the rate of increase of Patient's suffering as much as possible. How could this be a bad thing to do? Perhaps the problem is that we have given only an incomplete description of the states of affairs that constitute the consequences of Agent's action. It is true that Patient is suffering a good deal, but it is also true that he is not suffering even more. Perhaps this outweighs the suffering. More needs to be said, though. If I inflict some gratuitous suffering on you, it is true both that you suffer to a certain degree and that you don't suffer any more than that. This is so no matter how much suffering I inflict on you. But there's a crucial difference between *Doctor* and the infliction of gratuitous suffering. The suffering inflicted by Agent on Patient is not gratuitous. It is needed to prevent greater suffering. Perhaps what outweighs Patient's suffering is not the simple fact that he is not suffering even more, but the more complex fact that greater suffering is *prevented*. This latter fact un-

derlies our intuition that Agent has done a good thing. Patient was going to suffer even more, and Agent's action prevents that suffering. When I inflict gratuitous suffering on you, I don't prevent further suffering, because you weren't going to suffer even more. But now it seems we no longer have a noncomparative notion of consequences. What distinguishes the case of *Doctor* from the gratuitous infliction of suffering is that Patient was going to suffer even more if Agent hadn't administered the drug, but you weren't going to suffer even more if I hadn't inflicted suffering on you. If the fact that Patient doesn't suffer even more is to count as a consequence of Agent's action, but the fact that you don't suffer even more is not to count as a consequence of my action, our notion of consequences must involve comparisons with other possible worlds.

There is one more approach I will consider in search of a non-comparative notion of the goodness or badness of actions.¹⁶ If we can identify particular concrete states of affairs as the consequences of an action, we can evaluate those states of affairs as either on balance good, on balance bad, or neither. What prevents some (maybe most) states of affairs from being consequences of any particular previous action is that the action doesn't affect them. In *Doctor*, Agent's action affected Patient's conscious mental states, but it didn't affect the pleasures or pains of people thousands of miles away. The world that resulted from Agent's action contained Patient's suffering. It also contained the suffering and the pleasures of millions of other people, but, with respect to Agent's action, these states of affairs were unavoidable.¹⁷ Agent's action did not affect them. Consider an action A performed by an agent. A state of affairs S is avoidable with respect to A iff there is some action B that the agent could have performed in place of A, such that B would not have been followed by S.¹⁸ This suggests the following account of good actions:

¹⁶I owe the following suggestion to Peter Vallentyne.

¹⁷It would be more accurate to say that the states of affairs are either avoidable or unavoidable with respect to an agent at a time, since the crucial question is whether the agent could have acted at that time in such a way that the state of affairs would not have obtained. It makes for a less cumbersome formulation, however, to tie avoidability to the action.

¹⁸I mean here to exclude the use of backtracking counterfactuals. Only temporally subsequent states of affairs can be avoidable with respect to any particular action.

GS An action *A* is good iff the states of affairs that are avoidable with respect to *A* are, on balance, good; *A* is bad iff the states of affairs that are avoidable with respect to *A* are, on balance, bad.

Is *GS* an acceptable consequentialist account of good and bad actions? Is it a noncomparative account? It appears to be noncomparative. To determine whether a particular action is good, a consequentialist simply has to evaluate a set of states of affairs in one world. No comparison of values across worlds is needed. There is, however, an element of comparison across worlds involved in determining which states of affairs are avoidable with respect to an action. What makes a state of affairs avoidable with respect to an action is the fact that there is a possible world, accessible to the agent at the time of the action, that doesn't contain that state of affairs. I'm not sure whether this challenges *GS*'s claim to be a noncomparative account of good and bad actions. It isn't worth pursuing the point here, since *GS* is unacceptable, as I will explain.

Consider the following example, *Party*: Agent is at a party with one hundred other guests. The party is very isolated, so Agent is not in a position to affect the welfare of anyone who is not there. The other guests are all having a wonderful time. Agent possesses one dose of stomachache powder, which she drops in the glass of another guest, Patient, when Patient is not looking. Patient develops a severe stomachache as a result of ingesting the powder and is very unhappy for the rest of the party. The other ninety-nine guests, unaware of Patient's suffering, are extremely happy. Agent chose Patient because Patient was the most susceptible to the powder. She could have dropped the powder in the glass of any of the guests. Different guests would have suffered to different degrees, but none as much as Patient. She could also have disposed of the powder without harming anyone. If we have an account of good and bad actions, it must judge Agent's action to be bad, though perhaps not fiendishly so. *GS*, though, judges Agent's action to be good. Since Agent could have dropped the powder in anyone's glass, each guest's happiness or misery is an avoidable state of affairs with respect to Agent's action. Since ninety-nine of them are exceedingly happy and only one is very unhappy, the avoidable states of affairs are, on balance, good. (If you are inclined to think that the misery of one guest would outweigh the happiness of nine-

ty-nine, simply add more happy guests until you change your mind.) *GS* also judges many intuitively good actions to be bad. Agent's action in *Doctor* is one example.

I know of no other account of what it is for an action's consequences to be, on balance, good, or for the act to bring about more good than harm, or for its tendency to augment the general happiness to be greater than any it has to diminish it, or of any of the other related notions. Perhaps a satisfactory account can be produced, but I doubt it. Despite the widespread use of such notions, I conclude that they are, in fact, unavailable to consequentialists as accounts of the difference between good and bad actions.

5. Intuitive Judgments

Where does this leave a consequentialist account of the moral status of actions? Nothing I have said here challenges the traditional consequentialist account of rightness, though, as I said earlier, there are other reasons to do so. Consequentialists can judge actions to be better or worse than alternatives, and better or worse by certain amounts, but not to be good or bad simpliciter.

Just how surprising is this result? Perhaps it is not surprising at all. After all, moralities in which ascriptions of goodness to acts are fundamental are concerned with intentions or moral character in a way that consequentialism is not.¹⁹ Perhaps, then, it is not surprising that consequentialism can provide no satisfactory account of the goodness of actions. But I have been arguing for this conclusion by considering whether consequentialism can provide a satisfactory account of what it is for an action's consequences to be, on balance, good, or for the act to bring about more good than harm, or any of the other related notions with which I began the paper. That these notions, which are widely used by consequentialists, admit of several different interpretations, none of which provides a satisfactory consequentialist account of good actions, is certainly surprising. In fact, there are more surprises. Recall Sydney Carton's thoughts in *A Tale of Two Cities*. "It is a far, far better thing that I do, than I have ever done." He might plausibly have added, "And it's a pretty damn good thing, too." I have been arguing that

¹⁹I owe this point to an anonymous referee for the *Philosophical Review*.

the latter claim is, strictly speaking, unavailable to a consequentialist. Despite appearances, Carton's action is neither good simpliciter nor good to a certain degree (moderately, fairly, very, pretty damn, etc.). In fact, his former claim is also unavailable. Consequentialists can judge actions to be better or worse than alternatives, not better or worse than other actions performed at different times or by different people.

Why does this last result follow? Intuitively, a consequentialist should say that one actual action is better than another just in case the former produces a greater balance of goodness over badness (or smaller balance of badness over goodness) than does the latter. I have argued, however, that the notion of producing a balance of goodness over badness admits of several different interpretations, none of which provides the consequentialist with a plausible account of the goodness of actions. If Carton's action is neither good simpliciter nor good to a certain degree, what would be the basis for comparison with any of his past actions or anyone else's actions? In that case, it might be objected, what is the basis for comparison with any of his possible alternative actions? I do, after all, claim that an action can be assessed as better or worse than possible alternatives. Consider Carton's action in *A Tale of Two Cities* compared with the possible alternative of revealing his identity. His actual action is better than the alternative, just in case the world that results from it is better than the world that would have resulted from the alternative. If, as seems plausible, the world in which Carton is guillotined and Evremonde lives is better overall than the world in which Carton goes free and Evremonde is guillotined, then Carton's action is better than the alternative. How much better it is depends on how much better is the world that results from it. So, why not apply this technique to the comparison of different actual actions? In fact, it might even be easier to perform the comparison in this context. When we compare an actual action with a possible alternative, we are comparing the world that results from the action, the actual world, with the possible world that would have resulted from the alternative. For all the difficulty we have in assessing the actual world—given, for example, the difficulty of predicting the future²⁰—we have an even harder time assessing a

²⁰For an argument that the difficulty of predicting the future should not unduly worry consequentialists, see Norcross 1990.

merely possible world. When it comes to comparing different actual actions, however, we are not called on to assess merely possible worlds. Every actual action leads to an actual world, *the* actual world in fact. But this is precisely why the technique of comparing worlds will not give an acceptable means for comparing different actual actions. Every action leads to a temporal segment of the *same* world. Consider two actions, one performed before the other. The only difference between the world that results from the earlier action and the world that results from the later one is that the former world includes a temporal segment not included in the latter. If this segment is overall good, the former world is better than the latter; if the segment is overall bad, the former world is worse than the latter.²¹ If we apply the technique of comparing worlds to a comparison of different actual actions, we get the following unacceptable result: consider a temporal segment of the actual world, bounded by times t_1 and t_2 , such that the segment is overall good. Any action performed at t_1 is better than any action performed at t_2 . Worse still, any two simultaneous actions are equally good. Clearly, then, we can't compare different actual actions in the same way that we compare possible alternatives. Carton should have said, "It is a far far better thing that I do than anything I could have done instead."

I am not claiming that there are no comparisons for a consequentialist to make between different actions. Given the possibility of comparing an actual action with a possible alternative, a consequentialist can construct methods for comparing two different actual actions. She can compare them with respect to their distances from the best alternative in each case, or the worst alternative, or some other alternative. If it seems intuitively obvious that one action is better than another, it will probably also be obvious which comparison grounds that judgment. Different contexts will make different comparisons appropriate. In some contexts, perhaps no comparison will be appropriate.

Let me illustrate this point. First, consider an assessment of two different actions in which the relevant comparison is with the best

²¹I am assuming here that the morally relevant future is not infinite. For a discussion of possible problems for consequentialism if the morally relevant future is infinite, see Nelson 1991. For a suggested solution to these problems, see Vallentyne 1993.

alternative in each case. Recall the scenario in *Button Pusher*. Agent can push any one of ten buttons, killing between none and nine people, or push no button at all, with the result that ten people die. Suppose she pushes '5'. Now consider a variation. Agent is faced with only four buttons, labeled '1' to '4'. If she pushes no button, five people die. Suppose she pushes '2'. Intuitively, the former action is worse than the latter. Here, the relevant comparison in each case seems to be with the best alternative. The difference between pushing '5' and the best alternative of pushing '0' is greater than the difference between pushing '2' and the best alternative of pushing '1'. Pushing '2' is also closer to the *worst* alternative than is pushing '5', but that doesn't play a part in our judgment.

Next, consider a case in which the relevant comparison is with the worst alternative in each case. Suppose there are ten people trapped in a burning building. Agent can rescue them one at a time. Each trip into the building to rescue one person involves a considerable amount of effort, risk, and unpleasantness. It is possible, albeit difficult and risky, for Agent to rescue all ten. Suppose she rescues a total of five. Now imagine a similar situation, except that there are twenty people trapped in the building. Once again, Agent can rescue them one at a time. Each trip into the building involves the same amount of effort, risk, and unpleasantness as each trip in the last example. It is possible for Agent to rescue all twenty, though this would be even harder and more risky than rescuing ten. Suppose she rescues a total of seven. Intuitively, the rescue of seven in the latter case is better than the rescue of five in the former case. The relevant comparison here seems to be with the worst alternative in each case. The difference between rescuing seven and the worst alternative of rescuing none is greater than the difference between rescuing five and the worst alternative of rescuing none. Rescuing seven is also further from the *best* alternative of rescuing twenty than rescuing five is from rescuing ten, but that comparison isn't relevant to our intuitive judgment.

Finally, consider a case in which the relevant comparison is with neither the best nor the worst alternative in each case. Suppose that Ross Perot gives \$1000 to help the homeless in Dallas and I give \$100. Intuitively, Perot's action has better consequences than mine.²² Perot's action is further than mine from both the best al-

²²This is not to say that Perot's action is more praiseworthy or that it

ternative and the worst. It is plausible to assume that his immense riches give him a much greater range of options than is available to me. If we think that his action has better consequences than mine, we are not swayed by the fact that it falls short of his best alternative by more than does mine. But neither are we influenced by the fact that Perot's action is better than his worst option by a much greater amount than mine is better than my worst option. After all, if he gave only \$50, the gap between that and his worst alternative would still be far greater than between my giving \$100 and my worst option, and yet we would now judge my action to be better. So, why does Perot's action seem to have better consequences than mine? The natural comparison in this case seems to be with the alternative in which we do nothing with the money. The increase in goodness between the world in which Perot's money sits in the bank and the world in which he makes the donation is greater than the increase between the corresponding worlds involving my money.

6. Goodness and Context

Perhaps we can adapt the approach of the last section to provide an account of good and bad actions. If a consequentialist can explain why certain comparisons between actions and possible alternatives yield appropriate comparisons between different actual actions, why not use the same technique to give an account of good and bad actions? For example, if the appropriate comparison to invoke in comparing my donation with Perot's involves the alternatives in which we do nothing with the money, why not invoke the same comparison to determine whether each action is good or bad simpliciter? It does seem appropriate to judge both my donation and Perot's to be good actions. Isn't this because we judge each action to be better than the alternative in which the money simply sits in the bank? I have argued that a consequentialist can

shows him to be a better person. \$1000 is nothing to Perot, whereas \$100 is a significant amount to me. This consideration seems to affect the praiseworthiness of the action, not the goodness. Perot's action might not be more praiseworthy than mine, but isn't it still better? Rich people are simply better placed to do good than poor people are. This is not a reason to praise them or their actions (except inasmuch as such praise will encourage them to do more), or to denigrate poor people or their actions.

provide no satisfactory general description of the appropriate counterfactual comparison with which to judge an action's goodness. But perhaps a consequentialist doesn't need a *general* description of the counterfactual. Why not simply say that different comparisons will be relevant for judging different actions? This gives the following account:

GAP An action is good iff it is better than whichever possible alternative provides the appropriate comparison.

There is no general formula for identifying the appropriate comparison, but in most cases the context will make it clear.

I do not have a knock-down argument against *GAP*, but there are some fairly compelling reasons why a consequentialist should reject it as an account of the goodness of actions. First, part of the appeal of consequentialism is its simplicity and generality. *GAP* violates the spirit of consequentialism by using the notion of appropriateness without a general account of what makes a comparison appropriate.

Second, *GAP* allows nonconsequentialist factors, such as self-sacrifice, ownership, rights, and institutional duties to influence the question of whether an action is good or bad. Let me explain this point. In both the button-pushing and the burning building scenarios Agent's action determines how many people live or die. In the button-pushing scenario, anything other than the best action is bad. In the burning building scenario, anything other than the worst action is good. Our intuitions in these two cases seem to react to the elements of difficulty and self-sacrifice required of the agents. In *Button pusher*, the best action is no more difficult and requires no more self-sacrifice than any other action. In the burning building scenario, anything other than the worst action is very difficult and requires a significant degree of self-sacrifice. Similarly, consider a variation on Perot's donation to charity. Suppose that it is Agent's job to administer an endowment the size of Ross Perot's personal fortune. The purpose of the endowment is to help the homeless in Dallas. Agent is supposed to administer the endowment, to the best of her ability, to fulfill its purpose. Suppose that Agent gives \$1000, from the endowment of over \$2 billion, to help the homeless in Dallas. How does *GAP* assess this action? Given that it is Agent's *job* to help the homeless, the relevant compar-

ison is not with the alternative in which the money just sits in the bank. In Perot's case, however, this is the relevant comparison, because the money is his. Ownership gives him the right to do with the money whatever he wants (roughly). If, on the other hand, Agent just lets the endowment sit in the bank, she hasn't done her job.

The example of Agent administering the endowment also illustrates a third problem with *GAP*. Sometimes there is no clearly appropriate comparison. It may be obvious that Agent doesn't do a good thing by simply letting the money sit unused, but it is much harder to say just what she has to do with the endowment to do a good job. I suspect that many actions share this feature. Most of my examples in this paper have been designed to elicit strong intuitive reactions, but in many cases we just don't have strong intuitions about whether a particular action is good.

Perhaps the most serious problem with *GAP* is that it makes the goodness of actions context-relative. Consider again Perot's donation. Let's add a couple of details to the case: (i) Perot has a firm policy of donating up to, but no more than, \$1000 per month to charity. (Some months he gives less than \$1000, even as little as nothing at all, but he never gives more than \$1000.) (ii) He had been intending to give \$1000 this month to complete construction of a dam to provide water for a drought-stricken village in Somalia. As a result of Perot's switching the money this month to the homeless in Dallas, the dam takes another month to complete, during which time twenty children die of dehydration. Now it is not nearly so clear that Perot's action was good. A change in the description of the action might change the appropriate comparison. The extra details about the dam in Somalia make it unclear how to evaluate the action. It is still true that giving the \$1000 to the homeless is better than leaving it in the bank, but it is unclear whether this continues to ground the judgment that Perot's action is good. In fact, it is very tempting to say that Perot did a bad thing by diverting the money from the dam to the homeless. The problem here is not just that *learning* the details of the dam in Somalia changes the appropriate comparison. We could interpret *GAP* as incorporating a notion of appropriateness that assumes complete, or reasonably complete, knowledge. The problem is rather that what comparisons are appropriate can change with a change in the linguistic context, even if there is no epistemic change. For example,

different descriptions of the same action can make different comparisons appropriate. If we ask whether Perot's *diversion* of the \$1000 from the starving Somalians to the Dallas homeless was good, we will probably compare the results of the actual donation with the alternative donation to the Somalians. If, however, we ask whether Perot's *donation* to the Dallas homeless was good, we may simply compare the donation to the alternative in which the money sits in the bank, even if we know that Perot had previously intended to send the money to Somalia. Perhaps we'll say that the action was good, but not as good as the alternative of aiding the Somalians.

It might be objected at this point that there are theories of action individuation according to which Perot's *diversion* of the \$1000 from the starving Somalians to the Dallas homeless is not the same action as Perot's *donation* to the Dallas homeless. According to such theories, my example involves a switch from one action to another (spatiotemporally coextensive) one, rather than a mere switch in the way of describing a single action.²³ I do not know whether it is any less counterintuitive that two spatiotemporally coextensive actions should differ with respect to goodness than that the goodness of an action should depend on the linguistic context in which it is discussed. I do not need to settle the matter, however, since there can be changes in linguistic context that affect the appropriateness of comparisons without, on any plausible theory of action individuation, affecting which action is being referred to. There may be a change in the appropriate comparison even without a change of action description. Suppose that just before asking whether Perot's donation to the Dallas homeless was good, we have been discussing his prior intention to give the money to the Somalians. In this context, we are quite likely to compare the actual donation with the better alternative. On the other hand, suppose that just before asking whether his donation was good, we have been discussing the fact that Perot has made no charitable contributions at all in four of the last six months, and small ones in the other two. In this context, we will probably compare the actual donation with a worse alternative. No plausible theory of action individuation holds that the referent of an action term varies simply with what was

²³I am grateful to a referee for the *Philosophical Review* for bringing this objection to my attention.

being discussed at an earlier point in the conversation. In my example, however, the broad linguistic context in which the question is asked partly determines which comparison is appropriate and, according to *GAP*, whether a particular action is good or bad. Even without the first three arguments I brought against *GAP*, this should be enough to reject it as a theory of the goodness and badness of actions. The goodness of an action doesn't change with a change in the context in which it is discussed. It is much less implausible to suppose that the appropriateness of describing an action as good or bad changes with a change in linguistic context. If *GAP*, or something like it, is interpreted as a theory about what descriptions of actions are appropriate, it is quite consistent with my claim that a consequentialist should deny that any actions are good or bad simpliciter.²⁴

²⁴Perhaps there is a way to construct a consequentialist theory of good action in which different considerations determine whether different actions are good or bad, without accepting that the goodness of an action can change with the context in which it is discussed. Consider the following sketch of an account, suggested by a referee for the *Philosophical Review*. In any choice situation, the options available to an agent can be placed in (at least) a partial ordering, from best to worst, depending on the values of the worlds that result from them. An action is good (or bad) iff it's high (or low) enough in the partial ordering. The principles that determine the "enough" level vary from action to action. For example, if all but one of the alternatives involve causing gratuitous harm, only that one is high enough to be good, and the others are all low enough to be bad. If all but one involve conferring benefits, all but that one are high enough to be good. This would seem to be an account that mixes consequentialist and nonconsequentialist considerations. The "enough" level itself is not determined by the consequentialist value of the alternatives, but the partial ordering which the "enough" level cuts is. That might be reason enough for a consequentialist to reject this account. But some consequentialists (at least some of those who accept the arguments of this paper) might reason that a partially consequentialist account of good actions is better than no consequentialist account at all. However, on closer examination, this account reveals itself to be *wholly* nonconsequentialist. Consider one of the nonconsequentialist principles suggested above:

Harm If all but one of the alternatives involve causing gratuitous harm, only that one is high enough to be good, and the others are all low enough to be bad.

Why should we think that *Harm* will cut the partial ordering in such a way that all alternatives above the cut are good and all those below are bad? Isn't it possible that at least one of the alternatives that involves causing gratuitous harm will be ranked above the one that doesn't in the partial

7. Conclusion

I have argued in this paper that a consequentialist cannot give a satisfactory account of the goodness of actions in terms of the goodness of their consequences. My arguments also affect a consequentialist attempt to equate goodness (of actions) with rightness, or to give an account of the goodness of actions in terms of the value of the motives from which they spring. I don't claim, in this paper, that a satisfactory consequentialist account of good actions along one of these two latter lines *could* not be developed, though I strongly suspect that is the case. If my suspicion is correct, and, further, if the reasons for a consequentialist to reject rightness prove compelling, what *can* a consequentialist say about the moral status of actions? It appears that she can truly say of an action only how much better or worse it is than other possible alternatives. But common sense tells us that actions are (at least sometimes) right or wrong, good or bad. We look to moral theories to give us an account of what makes actions right or wrong, and good or bad. Indeed, it might even be thought that 'right action' and 'good action' are basic and indispensable moral concepts. How, then, can consequentialism claim to be a moral theory, if it can give no clear sense to such notions? This is an important objection, one that deserves more of an answer than I can give here. To the extent that you are moved by this objection, you may take the argument of this paper to be a contribution to the vast body of anticonsequentialist literature. Let me briefly explain why I don't so take it (leaving a detailed discussion of this issue for another paper).

It is true that 'right action' and 'good action' are concepts central to modern moral philosophy.²⁵ But are they indispensable, or

ordering? Indeed, isn't it precisely because of this possibility that we think of *Harm* as a *nonconsequentialist* principle? The same point applies to the principle about conferring benefits, and to any other nonconsequentialist principle that might be suggested. If we employ such principles to determine the goodness of actions, the consequentialist partial ordering is playing no role. Of course, we could specify that *Harm* only affects the goodness of actions in those cases in which no actions that involve causing gratuitous harm are ranked above any that don't in the partial ordering (and make similar specifications for all the other nonconsequentialist principles). But not only would this reduce the scope of the theory of good action considerably, it would also seem to be rather ad hoc.

²⁵Ancient moral philosophy, on the other hand, is centered around questions of virtuous character and the good life.

do they just appear to be so, because pretty much every competing theory, consequentialist and nonconsequentialist alike, gives some account of them? One approach to this question is to ask what the central function of morality is and how the concepts of ‘right action’ and ‘good action’ relate to it. I suggest, in common with many other moral theorists, that the central function of morality is to guide action by supplying reasons that apply equally to all agents.²⁶ The judgments that certain actions are right or good might seem to supply such reasons, but they are by no means the only sources of action-guiding reasons. Consequentialism, on my account, can tell an agent how her various options compare with each other. Concerning a choice between A, B, and C, consequentialism can tell the agent, for example, that A is better than B, and by how much, and that B is better than C, and by how much. She is thus provided with moral reasons for choosing A over B and B over C, the strength of the reasons depending on how much better A is than B and B than C. Doctors seem to provide prudential reasons of just this scalar nature when they tell us, for example, that the less saturated fat we eat the better.²⁷

“Nonetheless,” you might say, “even if consequentialism, on your account, survives as a moral theory, aren’t its chances of being true diminished by your arguments? Moral common sense tells us that some actions are good or bad, and better or worse than other actual actions. You say that consequentialism can make no sense of such judgments. To that extent, consequentialism clashes with moral common sense and should be rejected.” It is true that my version of consequentialism clashes with moral common sense in its denial that actions are, strictly speaking, good or bad, or better or worse than other actual actions. To the extent that you are immovably wedded to moral common sense, my arguments should give you reasons to reject consequentialism. However, those who are immovably wedded to the verdicts of common sense morality

²⁶Thus, if it is morally right to do x, all agents for whom x is an option have a reason to do x. The claim that *morality* provides reasons for action is different from the claim that particular *moral codes*, accepted by particular societies, provide reasons for action. Such codes may well provide reasons for action, but these reasons apply only to members of the relevant societies. The reasons provided by morality apply equally to all agents.

²⁷For a more detailed discussion of the action-guiding nature of scalar morality, see Howard-Snyder and Norcross 1993, 119–23.

do not need my arguments to reject consequentialism. Consequentialists are prepared to accept a considerable amount of disagreement with commonsense morality. I am not claiming, however, that consequentialists should simply ignore the deliverances of commonsense morality. At the very least, it is desirable for consequentialists to explain why it sometimes seems appropriate to make (and even express) the judgments that I have argued are strictly unavailable to them. That is why I have also argued that a consequentialist can explain the appropriateness of making judgments such as that the action was good, or better than a previous action, or worse than a different action performed by somebody else. I don't wish to make too much of this latter claim, though. Linguistic and moral appropriateness are two entirely different matters. To explain why a particular judgment sounds appropriate is not to justify the practice of making such judgments. One of the great strengths of consequentialist ethical theories is that they explain and justify some of our intuitions while challenging others. I suspect that there is *some* value in making the sorts of judgments about actions that I have argued are strictly unavailable to a consequentialist, but that overall we would be better off if we concentrated our attention on how our actions compare with other possible alternatives. When we contemplate a course of action, instead of asking whether we are doing better or worse than other people, we should ask whether there are better alternatives that we are willing to undertake. By focusing on each situation of choice, we may be less likely to become disheartened by what appears to be our inability to do as much good as others or complacent at our ability to do more. At the very least, we won't be tempted to justify our behavior by pointing out that many others are doing no better.²⁸

Southern Methodist University

²⁸While writing this paper, I encountered an early version of Christopher Hitchcock's paper "Farewell to Binary Causation" (1996). Hitchcock argues that "causal claims do not describe binary relations between causes and effects; instead, all such claims are made . . . relative to an alternative cause" (282). Some of Hitchcock's arguments are interestingly similar to my own, which leads us both to suspect that there may be a deep connection between the relativities of goodness and of causation. An exploration of this possible connection must await another occasion.

References

- Bentham, Jeremy. 1789. *Introduction to the Principles of Morals and Legislation*. Widely reprinted.
- Hitchcock, Christopher. 1996. "Farewell to Binary Causation." *Canadian Journal of Philosophy* 26:267–82.
- Howard-Snyder, Frances, and Alastair Norcross. 1993. "A Consequentialist Case for Rejecting the Right." *Journal of Philosophical Research* 18:109–25.
- . 1994. "The Heart of Consequentialism." *Philosophical Studies* 76: 107–29.
- Kagan, Shelly. 1989. *The Limits of Morality*. Oxford: Oxford University Press.
- Lewis, David. 1973. *Counterfactuals*. Cambridge: Harvard University Press.
- Mill, J. S. 1861. *Utilitarianism*. Widely reprinted.
- Nelson, Mark. 1991. "Utilitarian Eschatology." *American Philosophical Quarterly* 28:339–47.
- Norcross, Alastair. 1990. "Consequentialism and the Unforeseeable Future." *Analysis* 50:253–56.
- Quinton, Anthony. 1989. *Utilitarian Ethics*. 2d ed. La Salle: Open Court.
- Ross, W. D. 1973. *The Right and the Good*. Oxford: Oxford University Press.
- Scheffler, Samuel. 1982. *The Rejection of Consequentialism: A Philosophical Investigation of the Considerations Underlying Rival Moral Conceptions*. Oxford: Oxford University Press, Clarendon Press.
- Sidgwick, Henry. 1981. *The Methods of Ethics*. 7th ed. Indianapolis: Hackett Publishing Company.
- Slote, Michael. 1985a. *Common-Sense Morality and Consequentialism*. Boston: Routledge and Kegan Paul.
- . 1985b. "Utilitarianism, Moral Dilemmas, and Moral Cost." *American Philosophical Quarterly* 22:161–68.
- Smart, J. J. C. 1973. "An Outline of a System of Utilitarian Ethics." In J. J. C. Smart and Bernard Williams, *Utilitarianism For and Against* (Cambridge: Cambridge University Press).
- Vallentyne, Peter. 1993. "Utilitarianism and Infinite Utility." *Australasian Journal of Philosophy* 71:212–17.