Contextualism for Consequentialists

Alastair Norcross

Rice University, Houston, TX, USA

If, as I have argued elsewhere, consequentialism is not fundamentally concerned with such staples of moral theory as rightness, duty, obligation, moral requirements, goodness (as applied to actions), and harm, what, if anything, does it have to say about such notions? While such notions have no part to play at the deepest level of the theory, they may nonetheless be of practical significance. By way of explanation I provide a linguistic contextualist account of these notions. A contextualist approach to all these notions makes room for them in ordinary moral discourse, but it also illustrates why there is no room for them at the level of fundamental moral theory. If the truth value of a judgment that an action is right or good varies according to the context in which it is made, then rightness or goodness can no more be properties of actions themselves than thisness or hereness can be properties of things or locations themselves.

Keywords: contextualism, consequentialism, harm, right, good, linguistic.

1. Introduction

I have argued elsewhere[1] that consequentialism is not fundamentally concerned with such staples of moral theory as rightness, duty, permissibility, obligation, moral requirements, goodness (as applied to actions), and harm. In fact, I have argued that the standard consequentialist accounts of these notions are either indeterminate (in the case of the latter two) or redundant. What is fundamental to a consequentialist ethical theory is a value theory, for example hedonism or some other form of welfarism, and the claim that the objects of moral evaluation, such as actions, characters, institutions, etc. are compared with possible alternatives in terms of their comparative contribution to the good. For example, one action is better than another, just in case, and to the extent that, the world that contains it is better than the world that contains the other from the time of the choice onwards.

This assumes determinism, for the sake of simplicity. If indeterminism is true, we will have to replace talk of *the* world containing an action with talk of a set of

worlds. The relevant comparisons will involve something like the following account: 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 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.

Furthermore, our (moral) reasons for choosing between alternative actions, institutions, etc. are essentially comparative, and correspond to the comparative consequential value of the options. I might have a better reason for choosing to do A than to do B, and better by a certain amount, but neither reason is either good or bad *simpliciter*. So, if all a consequentialist moral theory supports at the fundamental level are comparative evaluations of actions, characters, institutions (and thus also comparative reasons for choosing among them), what, if anything, does it have to say about such notions as right and wrong, duty, obligation, good and bad actions or harm?

There are two main options, one of which I will be discussing. The first, which I will only briefly mention here, is a form of eliminativism, combined with an error theory regarding our common usage of these terms. The consequentialist could simply say that there's no such thing as right and wrong actions, good and bad actions, harmful actions, etc. It doesn't, of course, follow from this that "anything goes", if that is taken to mean that everything is permissible, and so, for example, it's perfectly permissible to torture innocent children. Just as no actions are either right or wrong, none are permissible or impermissible either. Neither does it follow that anything goes, if that is taken to mean that morality has nothing to say about actions. The action of torturing an innocent child will almost certainly be much worse than many easily available alternatives, and thus strongly opposed by moral reasons when compared with other options. It does, however, follow that descriptions of actions (or characters, or institutions) as being right or wrong, good or bad, harmful, required, permissible, and the like are all mistaken (either false or meaningless).

This might seem to be a rather uncomfortable result. We can understand how some, perhaps many, claims about the rightness or goodness or permissibility of actions are mistaken, but *all claims*? Is it plausible that we have *all* been mistaken all this time? I don't find this possibility particularly implausible. Similar things may well be true for certain areas of theological or scientific discourse. If there is no god, for example, all claims about what god loves or hates are mistaken (either false or meaningless). Similarly, much scientific discourse assumes the existence of entities that may turn out not to exist. It might, perhaps, be argued that the situation is different for morality. While theology and fundamental physics is unashamedly concerned with unobservable, or at least difficult to observe, entities, morality is concerned with everyday properties that require little or no expertise to discern. Although I don't find such considerations particularly compelling, I do

want to explore how a consequentialist can accommodate some of the commonly accepted moral properties, despite excluding them from the fundamental level of the theory.

What I propose is a form of contextualist analysis of the relevant moral terms, similar in form to some recent contextualist approaches to the epistemological notions of knowledge and justification. Roughly, to say that an action is right, obligatory, morally required, etc. is to say that it is at least as good as the appropriate alternative (which may be the action itself). Similarly, to say that an action is good is to say that it resulted in a better world than would have resulted had the appropriate alternative been performed. To say that an action harmed someone is to say that the action resulted in that person being worse off than they would have been had the appropriate alternative been performed. In each case, the context in which the judgment is made determines the appropriate ideal or alternative. I will illustrate first with the cases of good actions and harmful actions, and then say a little about a contextualist analysis of 'right'.

2. Problems with Non-Contextualist Accounts of Good and Harm

In order to explain (and motivate) a contextualist account of good actions and harmful actions, I will briefly explain why satisfactory noncontextualist accounts of such notions are not available to the consequentialist.[2] If the goodness of an action is to be a consequentialist property, something like the following account suggests itself:

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

The general idea expressed in **G** is used by philosophers, both consequentialist and non-consequentialist,[3] though not necessarily as an explicit account of good and bad actions. But what does it mean to produce more goodness than badness, or, to put it another way, to have consequences that are on balance good?

The obvious answer is that for an action to have on balance good consequences is for it to make a positive difference in the world, that is, to make the world better. But better than what? A first attempt is to say better than it was before the action. But this clearly won't do. To see this, consider an example in which the world contains only two sentient beings, Agent and Patient. 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. So the result of administering the drugs is that Patient's suffering con-

tinues to increase, but at a slower rate than he would have experienced without them. The very best thing she can do has the consequence that Patient's suffering increases. The world is worse after Agent's action than it was before, but Agent's action is clearly not on that account bad. In fact, inasmuch as a consequentialist is inclined to make a judgment about the action's goodness, she would say that it is good.

In evaluating actions, a consequentialist compares states of affairs, not across times, but across worlds. The reason why it seems as if Agent's action is good is that it does make the world better, not better than it was, but better than it *would have been* if the action hadn't been performed. This suggests the following account of good actions:

> **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 explains why Agent's action is good. If she hadn't administered those drugs to Patient, Patient would have suffered even more. But this is an easy case, which hides a crucial problem with GC. According to GC, whether an action is good or bad depends on what the world would have been like if it hadn't been performed. So, what would the world have been like, if Agent had administered those drugs to Patient? That depends on what Agent would have done instead. She might have tried a different course of treatment, which was less effective. She might have simply sat and watched while Patient's suffering increased. She might have tried a different course of treatment that actually increased the rate of increase of Patient's suffering (either intentionally or not). In this case, we don't need to know precisely what Agent would have done instead, because we know that she did the best she could, and thus that the world would have been worse, if she had done *anything* else.

But other examples are not so easy. Consider the following:

> *Button Pusher*. Agent can push any one of ten buttons (labeled '0'through '9'), killing between none and nine people, or push no button at all, with the result that ten people die. No button is any more difficult to push than any other, nor is there any pressure (either physical or psychological) exerted on Agent to push any particular button.

Suppose that Agent pushes the button labeled '9', with the result that nine people die. Intuitively, this seems like a pretty bad action. However, suppose also that Agent is highly misanthropic, and wants as many people as possible to die. Her initial inclination was to press no button at all, so that all ten would die. She also enjoys being personally involved in the misfortunes of others, however, and believes that pressing a button would involve killing, whereas refraining from pressing any button would involve 'merely' letting die, which, from her misguided perspective, is less personally involving. She struggled long and hard over her decision, weighing the advantage of one more death against the disadvantage of less personal involvement. She never

contemplated pressing any button other than '9'. It's clear, then, that if Agent hadn't pressed '9', she would have pressed no button at all. So the world would have been worse, if she hadn't pressed '9'. But this doesn't incline us to judge her action to be good.

Although *Button Pusher* might suggest that anything less than the best action is bad, we are not likely to endorse that as a general principle. Consider:

> *Burning Building.* 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 that Agent rescues nine people, and then stops, exhausted and burned. She could have rescued the tenth, so doesn't do the very best she can, but do we really want to say that her rescue of nine people wasn't good (was actually bad)?

None of the different 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 are 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.

Consider now the consequentialist approach to harm:

> HARM An act A harms a person P iff P is worse off, as a consequence of A, than she would have been if A hadn't been performed. An act A benefits a person P iff P is better off, as a consequence of A, than she would have been if A hadn't been performed.[4]

It is easy to see that the same problems that apply to the consequentialist account of good and bad actions apply to the consequentialist account of harmful and beneficial actions. The following example will illustrate: suppose you witness the following scene at Texas Tech University: A member of the Philosophy department, passing Bobby Knight on campus, waves cheerily and says "Hey, Knight." Bobby Knight, turning as red as his sweater, seizes the hapless philosopher around the neck and chokes her violently, while screaming obscenities. By the time Bobby Knight has been dragged away, the philosopher has suffered a partially crushed windpipe and sustained permanent damage to her voicebox, as a result of which she will forever sound like Harvey Fierstein.

Has Bobby Knight's act harmed the philosopher? The intuitive answer is obvious, and HARM seems to agree. The philosopher is much worse off than she would have been had Bobby Knight not choked her (unless, perhaps, she has always wanted to sound like Harvey Fierstein). But suppose we discover that Bobby Knight has recently been attending anger management classes. Furthermore, they have been highly successful in getting him to control his behavior. When he be-

comes enraged, he holds himself relatively in check. On this particular occasion (only the third violent outburst of the day), he tried, successfully, to tone down his behavior. In fact, if he hadn't been applying his anger management techniques, he wouldn't have choked the philosopher, but would rather have torn both her arms from her body and beaten her over the head with them. Since it took great effort on Bobby Knight's part to restrain himself as much as he did, it seems that the closest possible world in which he doesn't choke the philosopher is one in which she is even worse off. HARM, in this case, seems to give us the highly counterintuitive result that, not only does Bobby Knight's act of choking not harm the philosopher, it actually benefits her. HARM also seems to give the result in *Button Pusher* that Agent doesn't harm any of the nine people who die as a result of pushing '9'. They are no worse off than they would have been if she hadn't pressed '9'. If she hadn't pressed '9', she wouldn't have pressed any button, and all ten people would have died.

As with good and bad actions, the consequentialist account of harmful and beneficial actions includes a comparison with an alternative possible world. To harm someone is to make her worse off than she *would have been*. The alternative with which we are to compare the actual action, though, is not always plausibly identified by the counterfactual. Features of the conversational context in which a particular action is being assessed can affect which alternative is the appropriate one.[5]

3. Contextualism about Good, Harm and Right

Consider first a contextualist account of good action:

> **G-con** An action is good iff it is better than the appropriate alternative.

As examples for which the conversational context is unlikely to change the appropriate alternative, consider again *Button Pusher and Burning Building*. Suppose that Agent pushes '5' in *Button Pusher*. It is hard to imagine a conversational context in which anything other than pushing '0' is selected as the appropriate alternative. Pushing '5' would clearly be judged a bad action in just about any plausible conversational context. Now suppose that agent rescues three people in *Burning Building*. In most conversational contexts the appropriate alternative will be rescuing none (or perhaps one), and so the rescue of three will be judged to be good.

Now consider an example for which conversational context might change the appropriate alternative.

> *Perot*. Ross Perot gives $1000 to help the homeless in Dallas and I give $100.

In most conversational contexts both of our actions will be judged to be good, because the appropriate alternatives will be ones in which we give no money. But 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 on 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 we should say 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 point here is not just that *learning* the details of the dam in Somalia changes the appropriate comparison. The point 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 Somalis to the Dallas homeless was good, we will probably compare the results of the actual donation with the alternative donation to the Somalis. 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 Somalis.

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 Somalis 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. However, there can clearly be changes in linguistic context that affect the appropriateness of comparisons, without affecting which action is being referred to, on any plausible theory of action individuation. 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 Somalis. 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.

Now lets consider a contextualist account of harm:

> **H-con** An action A harms a person P iff it results in P being worse off than s/he would have been had the appropriate alternative been performed.

Many straightforward examples involve actions for which the conversational context is most unlikely to change the appropriate comparison, or at least unlikely to change it so as to produce a different judgment. For example, chancing to encounter you at a philosophy conference, I kill and eat you. It is hard to imagine a conversational context in which the appropriate alternative action is worse for you than being killed and eaten. Likewise, to use a real example, if I say that Booth's shot harmed Lincoln, the context selects, as an appropriate alternative act of Booth, pretty much anything else except shooting Lincoln. It may be true that Booth could have shot Lincoln in such a way as to lead to a much more agonizing death than the one he in fact suffered. This alternative, however, is normally not salient (and may never be). However, it's also a fairly straightforward matter to produce an example for which the appropriate alternative does change with the conversational context.

Sometimes, different, equally normal, contexts can render one act a harming or a benefiting. For example, my father writes a will, in which I receive half his estate. This is the first will he has written. Had he died intestate, I would have received all of his estate. Two among his many other options were to leave me none of his estate or all of it. Does my father's act of will-writing harm me or benefit me? Imagine a conversation focused on my previous plans to invest the whole estate, based on my expectation that I would receive the whole estate. It might be natural in such a context to describe my father's act as harming me. I end up worse off than if he had left me all his estate, which I had expected him to do, either by not making a will at all, or by making one in which he left me the whole shebang. Imagine, though, a different, but equally natural, conversation focusing on my lack of filial piety and the fact that I clearly deserve none of the estate. In this context it may be natural to describe my father's act as benefiting me. After all, he *should* have left me nothing, such a sorry excuse for a human being I was.

At this point an objection may arise. Introducing the previous example, I said that different contexts can render one act a harming or a benefiting. Given that I am talking about harm *all things considered*, how can I claim that one act can correctly be described as both a harming and a benefiting? Wouldn't this be contradictory? Likewise, in discussing Perot's donation to the Dallas homeless, I said that the context in which it is discussed can determine whether the appropriate comparison is with a better or a worse alternative, and thus whether Perot's action is correctly described as good or bad. Again, it seems that I am claiming that one action can be correctly described as both good and bad. Isn't this contradictory? No. In order to see why not, we need to be precise about what I am committed to. I say that one act can be correctly described *in one conversational context* as good, and can be correctly

described *in a different conversational context* as bad. The reason why no contradiction is involved is that a claim of the form 'act A was good' can express different propositions in different contexts. (The same point, of course, applies to the contextualist accounts of harm and other moral notions.) On my suggested account of good actions, to claim that act A was good is to claim that A resulted in a better world than would have resulted if *the appropriate alternative* to A had been performed. Given the context-relativity of *the appropriate alternative*, claims about good and bad actions have an indexical element. Just as 'today is a good day to die' can express different propositions in different contexts of utterance, so can 'Perot's donation to the Dallas homeless was good'.

At this point I should clarify the role of salience in my contextualist account of moral terms. I mean by salience, roughly, the degree to which the participants in a conversational context consciously focus on an alternative. There may be more sophisticated accounts of salience, but this is certainly a common one. Salience often plays a role in determining which alternative the context selects as the appropriate one, but salience may not be the only determining factor. To see this, consider an example that might be thought to pose a problem for my account, if salience is solely responsible for selecting the appropriate alternative.[6] Imagine a group of comic-book enthusiasts talking about how great it would be if their leader, Ben, had the abilities of Spiderman. After an hour or three of satisfying fantasizing, they are joined by Ben himself, who apologizes for being late. He explains that he was on his way when his grandmother called him on his cellphone. She had fallen, and she couldn't get up without his help. It took him more than an hour to get to her, because of traffic congestion, during which time she had been lying uncomfortably on the floor. Once he helped her up, though, she was fine. He is sorry that he is late, but the rest of the group, who are also devoted grandsons, must agree that benefiting his grandmother is a good excuse. "Au contraire", reply his friends, that is the "worst excuse ever". He didn't benefit his grandmother at all, but rather harmed her, since he would have reached her a lot sooner, and prevented much suffering, if he had simply used his super spider powers to swing from building to building, instead of inching his way in traffic. Furthermore, he would have reached the meeting on time. Clearly, something is amiss here. Even though the alternative in which Ben swings through the air on spidery filaments is, in *some* sense, salient, it is not thereby the appropriate alternative with which to compare his actual behavior. We can't make an alternative appropriate simply by talking about it, although we may be able to make it salient that way. Perhaps we should add to salience, among other things, a commitment to something like 'ought implies can'. Since Ben cannot swing through the air on spidery filaments, this is ruled out as an appropriate alternative.[7] I don't here have the time (or the inclination) to give a detailed account of how conversational context determines the appropriate alternative. I suspect that the correct account will be similar to the approach of contextualists in epistemology, such as David Lewis, Mark Heller, and Keith deRose.

Finally, consider a contextualist analysis of 'right':

>   **R-con.**   An action is right iff it is at least as good as the appropriate alternative.

The idea here is that the concept of right action (and duty, permissibility, obliga-
tion, and the like) invokes a standard, against which the action in question is
judged. The standard maximizing consequentialist theory is a non-contextualist
theory of the right, which fixes the standard as optimizing. For the maximizer, the
appropriate ideal is always the optimal option. However, the contextualist ap-
proach I am suggesting allows the conversational context to affect the standard. It
seems likely that most (ordinary) contexts will be sensitive to such factors as diffi-
culty (both physical and psychological), risk, and self-sacrifice in establishing the
appropriate ideal. For example, most, if not all, contexts will establish the act of
pushing '0' as the appropriate ideal in *Button Pusher*, so that any other action will
be judged wrong. *Burning Building*, is a little trickier, but it is hard to imagine
many ordinary contexts that set the rescue of *everyone* as the appropriate alterna-
tive. The standard criticism of maximizing consequentialism that it fails to accom-
modate supererogation is based in the intuition that there are cases in which duty,
or right action, doesn't demand maximizing. *Burning Building* seems to be a good
example of one. In order to get a context that would set the rescue of all ten people
as the appropriate ideal, we could imagine a conversation among committed max-
imizing consequentialists, or perhaps among proponents of a Christ-as-ideal moral
theory, or perhaps it will be enough to imagine a conversation in a philosophy class
that has just been presented with maximizing consequentialism. Just as the episte-
mological contextualist presents classroom contexts as setting particularly de-
manding epistemic standards, and thus as being ones in which "I don't know that I
have hands" can be uttered truly, so the ethical contextualist can present classroom
contexts in which the maximizing alternative determines the truth value of claims
of rightness. Of course, classroom contexts might also set very low standards. A
discussion of the demandingness objection to consequentialism might set a pretty
lax standard.

It is important to stress that the contextualism I am suggesting is a *linguistic*
thesis. I am not suggesting that the rightness (or goodness, etc.) of a particular ac-
tion can vary with the context in which it is discussed. I am suggesting that a sen-
tence such as "Michael Moore was morally right to describe Bush as a 'deserter'"
may express different propositions when uttered in different contexts. The right-
ness of Moore's act (of describing) doesn't vary with the context in which it is dis-
cussed. That is because the context in which the previous sentence was uttered (or
read) determined the property picked out by 'rightness' *in that context*. Assume
that Moore's act possessed *that* property. If so, no change in linguistic context can
change the fact that Moore's act possessed *that* property. A change in linguistic
context can make it the case that a different utterance of 'rightness' will pick out a
different property.

A contextualist approach to all these notions makes room for them in ordinary moral discourse, but it also illustrates why there is no room for them at the level of fundamental moral theory. If the truth value of a judgment that an action is right or good varies according to the context in which it is made, then rightness or goodness can no more be properties of actions themselves than thisness or hereness can be properties of things or locations themselves. To be more accurate, since 'right' (and the other terms I have discussed) can be used to pick out different properties when used in different contexts, many actions will possess *a* property that can be legitimately picked out by 'right' (or 'good', 'harmful', etc.) and lack many other such properties.

Notes

1.   "Good and Bad Actions", *The Philosophical Review*, Vol. 106, No. 1; January 1997, pp. 1–34; "A Consequentialist Case for Rejecting the Right", *The Journal of Philosophical Research*, Vol. 18; 1993, pp. 109–125, co-authored with Frances Howard-Snyder; "Reasons and Demands: Rethinking Rightness", in James Dreier (ed.) *Blackwell Contemporary Debates in Moral Theory,* forthcoming; "Harming in Context", *Philosophical Studies,* 2004 forthcoming; "Scalar Act-Utilitarianism", in Henry R. West (ed.) *Blackwell Guide to Mill's* Utilitarianism*, forthcoming.
2.   For fully detailed accounts, see "Good and Bad Actions", and "Harming in Context", Op. Cit.
3.   For specific examples, see "Good and Bad Actions" op. Cit., 5–7.
4.   Something like this principle is assumed in most consequentialist writing about harmful and beneficial actions. See, for example, Derek Parfit, *Reasons and Persons*.
5.   Counterfactuals themselves are, of course, infected by context, but not always in the same way as judgments about harm. For example, simply entertaining a counterfactual may change the context in a way in which considering a judgment of harm without explicitly entertaining the counterfactual may not.
6.   I owe at least the general idea of this example, though not the details, to Ben Bradley. He suggested something like this in discussion as a problem for my account.
7.   I owe this suggestion to Julia Driver.

Alastair Norcross
Department of Philosophy
Rice University
6100 Main Street, MS 14
Houston, TX 77005
USA
e-mail: norcross@rice.edu