Kadish Center for Morality, Law & Public Affairs

"Consequentialism, Time, and Value"

David Killoren
University of Wisconsin, Madison

Consequentialism, Time, Value, and Common-Sense Morality
David Killoren

Abstract

Is consequentialism consistent with common-sense morality? I argue for a negative answer to this question. In Sections 1-4, I develop and defend a definition for "consequentialism." In Section 5, I attempt to show that, given this definition, consequentialism and common-sense morality cannot be reconciled. In Section 6, I argue that, on the definition of consequentialism I defend, consequentialism should be understood, not as a view about the relationship between the *deontic* and the *evaluative* (as many philosophers suppose), but as a view about the relationship between the deontic and *time*.

Introduction

Definitions for "consequentialism" tend to be either *stipulative* or *descriptive*.[1] Stipulative definitions simply set forth a plan for usage of the term defined; they make no claim to represent the usage patterns of anyone but a current speaker in a current context. Descriptive definitions seek to identify the definition implicit in the usage patterns of competent members of a linguistic community in a specified range of contexts. Since "consequentialism" is a term of art – a bit of philosophers' jargon – it is commonly supposed that all useful or informative definitions for it must be merely stipulative.[2] This reflects an odd prejudice against terms of art. Philosophers have exhibited abiding interest in finding descriptive definitions for words in common use like "knowledge," "obligation," "being," and so on.[3] Presumably, the reason for such interest is that (a) philosophers care to know about the concepts such words invoke in standard usage – and (b) philosophers believe that descriptive definitions can (at least partially) reveal those concepts to us. Thus, if philosophers are not equally interested in finding descriptive

---

[1] There is also the possibility of *explicative* definition (c.f. Carnap (1956); for general discussion, see Gupta's useful (2008) encyclopedia entry). It is difficult to characterize explicative definition in a footnote, but the gist is that an explicatrive definition for a term is meant to remain close enough to common usage to capture the concept invoked by it, but to depart from common usage so far as is necessary to serve whatever local purposes the definition is provided to serve. Explicative definition, it is sometimes said, *recommends a conceptualization* of a concept already in common use. Explicative definition is thus supposed to provide a kind of middle ground between description and stipulation, and therefore to have at least some of the advantages of each. However, it seems to me that the activity that ordinarily goes by the name "explicative definition" often has *neither* of these two sets of advantages: In *departing* from common usage, explicative definition risks failing to capture the concept *invoked* in common usage; but in remaining *loosely tethered* to common usage, an explicative definition risks failing to serve the local purposes of the philosopher who devises it. In this way, in attempt to serve two masters, explicative definition may fail to serve any. In any case, serving one of these masters is trouble enough.

[2] This supposition, I grant, is more often articulated in conversation than in print. But Portmore mentions it in his (2007).

[3] Here it might be objected that what philosophers *really* show concern for, in discussing the meanings of these words, is *explicative* definition for these terms (see fn. 1). I doubt this. For one thing, I think that philosophers' reliance on the *intuitions* of ordinary, competent speakers in their arguments betrays an interest in descriptive definitions (because I think such intuitions are primarily useful as evidence regarding the best descriptive definition for a term). But there is not space to support this view here.

definitions for their own terms of art, then this must be because (c) philosophers do not care to know about the concepts such terms invoke in standard usage, or because (d) descriptive definitions cannot (even partially) reveal those concepts to us. But (c) is at least *puzzling*: Why *shouldn't* philosophers care to know about the concepts invoked by their own terms of art? Further, (d) seems likely to be false if (b) is true.

Moreover, it seems doubtful that there is a principled distinction between common-use terms and terms of art. Philosophers talk about different things than ordinary people talk about, so it stands to reason that some of philosophers' terms of art would lack a correlate in common use. But this is just a byproduct of the peculiarity of philosophical concerns; it *could* have been that the concepts for which philosophers have names would also have been given names by ordinary people. If this had been the case, would we suddenly find ourselves interested in finding the descriptive definitions for the common-use terms? Surely not. Thus, it is difficult to see how there can be a *principled* blanket dismissal of concern for descriptive definitions for philosophical jargon.

But even if this is right, we still have not been given a reason to think that the descriptive definition for "consequentialism," in particular, should interest us. Since much of this paper will be devoted to the quest for such a definition, I need to explain why the quest for such a definition is not boring.

It is an open, and purportedly interesting, question whether consequentialism is consistent with common-sense morality. Call this question Q1. Not only does Q1 *sound* interesting, but many philosophers have weighed in on it;[4] presumably, they would not have done so if it were a boring question. But Q1 *would* be boring if, in trying to answer it, we were meant to rely solely on *stipulative* definitions for "consequentialism." This is because it is very easy to stipulate a definition that guarantees (or makes impossible) consequentialism's compatibility with common sense. But suppose Q1 were, instead, equivalent to the following question:

> Q2: Is *the sort of moral theory which we (i.e. members of the philosophical community) have in mind when we talk about consequentialism* consistent with common-sense morality?

Q2, I submit, is obviously not boring (or at least not obviously boring). Moreover, it seems plain that Q2 is the question under consideration by at least some of those who discuss whether consequentialism is consistent with common-sense morality.[5] But to find an answer to Q2, we need a descriptive definition (or functionally adequate approximation thereof) for "consequentialism." Thus, it seems to me, we should take an interest in finding the descriptive definition for "consequentialism."

The main purpose of this paper is to show that the answer to Q2 is "No;" this is what I shall argue in Section 5. To defend this answer, I will need to develop and defend a descriptive definition for "consequentialism;" this will be my task in Sections 1-4.

1. Two candidate definitions for "consequentialism"

---

[4]   See, e.g. Foot (1983), Schroeder (2007), Portmore (2007).
[5]   This is especially apparent in Foot's (1983) discussion.

Sinnott-Armstrong (2006) has defined consequentialism as the view that "whether an act is morally right depends only on *consequences* (as opposed to the circumstances or the intrinsic nature of the act or anything that happens before the act)." There are a few ways in which we might object to this definition. We might point out, for instance, that the definition does not clarify that the consequences on which an act's rightness is supposed to depend are not only the act's *own* consequences, but also on the consequences of acts one might have performed instead. Or we might point out that the definition seems to assume that the consequences of an act are not part of its circumstances; it is not obvious that this is so, and it is seems odd to suggest that all consequentialists are, by virtue of their consequentialism, committed to say that it is so. Or we might point out that the definition says nothing about deontic statuses other than rightness. Does this mean that consequentialism has no implications regarding how we may determine wrongness, permissibility, supererogation, etc.?

But these are quibbles; set them to one side. A more interesting ground for objection is this: that *evaluative* notions play no role in the definition. Nowhere in this definition do words like "goodness" or "value" appear. Contrary to this, it is widely held that consequentialism has something *essentially* to do with the evaluative.[6] Of course, this is not a fact of which Sinnott-Armstrong is simply unaware; but he seems to think that evaluative notions are essential only to some *forms* of consequentialism. In particular, Sinnott-Armstrong defines "evaluative consequentialism" as the view that "moral rightness depends only on the *value* of the consequences (as opposed to other features of the consequences)."

Still, many would object to Sinnott-Armstrong's classification scheme.[7] They would object on the ground that, if one denies the view Sinnott-Armstrong calls "evaluative consequentialism," she thereby repudiates consequentialism itself. On this view, consequentialism, in standard (philosophical) usage, *just is* "evaluative consequentialism." In fact, this seems to me to be the standard view; it is extremely common for definitions of consequentialism to take a form equivalent, in the relevant respects, to the definition Sinnott-Armstrong provides for evaluative consequentialism (see fn. 6).

I am going to argue against this view; that is, I will argue that Sinnott-Armstrong's definition for consequentialism is, in its essentials, correct. But I will also argue that Sinnott-Armstrong is incorrect to classify evaluative consequentialism as a subtype of consequentialism. This is because, on the view I will defend, *some* (but not all) forms of evaluative consequentialism fail to be genuine forms of consequentialism.

It will prevent confusion (as well as some possible question-begging) if we rename Sinnott-Armstrong's two definitions for consequentialism. So let us simply call them D1 and D2, as follows:

---

[6] There is no shortage of philosophers to cite in support of this claim. For two authoritative examples, see Kagan (1989) and Darwall (2002); in both cases a definition for consequentialism essentially equivalent to Sinnott-Armstrong's "evaluative consequentialism" (see discussion below) is provided.

[7] In particular, those who advance definitions like Kagan's and Darwall's would object. However, it is important to note that Sinnott-Armstrong's style of definition is far from unprecedented. For instance, I think it is arguable that Sinnott-Armstrong's definition comes quite close to capturing the original *intended* meaning for consequentialism, i.e. the meaning intended in Anscombe (1958), wherein "consequentialism" was coined.

D1 = whether an act is morally right depends only on *consequences* (as opposed to the circumstances or the intrinsic nature of the act or anything that happens before the act).

D2 = moral rightness depends only on the *value* (or goodness – I will use these terms interchangeably) of the consequences (as opposed to other features of the consequences).

To repeat: On the view I will defend, D1 represents a correct descriptive definition for consequentialism; and D2 represents an incorrect definition for consequentialism: some moral theories affirm D2 while failing to be genuine forms of consequentialism.

2. D1 and D2 are not equivalent

Before defending D1 and criticizing D2, I want to show that D1 and D2 are not equivalent. It may seem that this step is unnecessary, but many philosophers proceed as though the two definitions are interchangeable,[8] so it is worth pausing to explain why this is not appropriate. We can do this in two steps: first, by showing that D2 does not imply D1; second, by showing that D1 does not imply D2.

D2 does not imply D1. Assume D1. Then the value of the consequences of one's actions (entirely) determines[9] whether they are right or wrong. If (i.) the consequences of one's actions (entirely) determine their own value, and (ii.) determination is transitive,[10] it would follow that (iii.) the consequences of one's actions (entirely) determine whether they are right or wrong. But (i.) is likely false, and there is no reason to suppose consequentialists are committed to it. E.g., suppose a prison guard executes a prisoner. Is the consequence of this action (i.e. the prisoner's death) good or bad? Arguably, whether it is good or bad is determined (at least in part) by whether the prisoner was a violent criminal or an innocent falsely accused. Yet the prisoner's criminal past (or lack thereof) is clearly not a consequence of the guard's action. Thus, the value of the consequences, in this case, is plausibly not *entirely* determined by the consequences themselves. See *Figure* below.

---

[8] For especially clear cases in which philosophers have proceeded in this way, consider Slote (1985) and Shaw (2006).

[9] In this paper, I am relying on an intuitive, contextual grasp of determination: Determination is the sort of relation that we have in mind when we assert that pleasure, or desire-satisfaction, or whatever, *makes* some outcome good – or when we assert that the badness (relative to alternatives) of an action's consequences *make* the action wrong; etc. There is not space here to defend an account of determination complete enough to decide whether determination really is transitive. Obviously, if determination is not transitive, this only bolsters my case here, but for the sake of argument I grant that it is transitive.

[10] Determination is transitive exactly if the following conditional is true (for any A, B and C):

[A determines B and B determines C] only if [A determines C].

D1 does not imply D2. Assume D1. Then the consequences of one's actions (entirely) determine whether they are right or wrong. Plainly, this doesn't imply that the *value* of the consequences of one's actions (entirely) determines whether they are right or wrong. See *Figure* below.

3. Direct and Indirect Utilitarianism

I have just argued that D1 and D2 are not equivalent. In this and the next section, I want to argue that D1 is superior, as a descriptive definition, to D2.

Consequentialism is a family of moral theories. A definition for consequentialism provides criteria by which to determine whether a given moral theory belongs to that family. The criteria provided by a (good) *descriptive* definition should match the usage patterns of competent speakers. Thus, if a given moral theory seems, to competent speakers, to be a version of consequentialism, then a candidate descriptive definition had better not classify it as a version of non-consequentialism (and vice-versa). Given this, and given that we take *ourselves* to be competent speakers, we should be able to rank the

adequacy of candidate definitions by comparing how well they match our "intuitions," i.e. our judgments as to whether various moral theories are consequentialist moral theories. (Of course, this is not a fool-proof plan. Competent speakers can misapply words, especially when they are considering unusual cases. I will return to this point below.)

Consider the following moral theory:

Indirect Utilitarianism: (a) In any given circumstance, one's sole obligation is to act so as to maximize the good (i.e. value – recall that we are using "goodness" and "value" interchangeably) and (b) only pleasure is good, and only the absence of pleasure is bad (and each same-sized unit of pleasure is equally good, regardless of who experiences the pleasure, when it is experienced, etc.)

Indirect Utilitarianism is, unquestionably, a consequentialist moral theory; indeed, as a version of Benthamite utilitarianism, it is arguably the very prototype of consequentialism. In observing this, I remain neutral between D1 and D2; Indirect Utilitarianism is a consequentialist theory on either definition.

Indirect Utilitarianism has many flaws. Here is one line of argument against it. Judith Thomson has argued (1994) that "there is no such thing as goodness." Simplifying somewhat, her view is that there is "goodness-in-a-way" (e.g. *being a good swimmer* or *being good for John but not for Sue,* etc.), but not simple *goodness* of the type that consequentialists normally have in mind. Suppose an Indirect Utilitarian encounters this view and is convinced by Thomson's argument for it. Then the Indirect Utilitarian would reject (b), and would thereby cease to be an Indirect Utilitarian. However, suppose that the (now former) Indirect Utilitarian is *not* convinced to give up the deontic upshot of her view. That is, this thinker continues to believe that one ought to act so as to maximize pleasure; only now she justifies this view in some way other than by reference to pleasure's *goodness*. So she adopts the following alternative view:

Direct Utilitarianism: (c) In any given circumstance, one's sole obligation is to act so as to maximize pleasure (without regard to who experiences the pleasure thus maximized, or when it is experienced, etc.); yet (d) nothing whatsoever, not even pleasure, is good.[11]

Now, it seems to me to be quite clear that, in shifting from Indirect Utilitarianism to Direct Utilitarianism, our thinker does not give up *consequentialism*. Direct Utilitarianism is, to be sure, an *odd* view, but it is hard to imagine a (competent, unbiased) user of philosophical jargon would not classify it as a consequentialist theory. If this is correct, then we cannot accept D2 as the basis for a definition for consequentialism. This is because Direct Utilitarianism is inconsistent with D2: Direct Utilitarianism denies that the rightness of one's actions depends on the value of the consequences. D1, however, as a definition for consequentialism, remains unscathed, since both Direct Utilitarianism and Indirect Utilitarianism imply D1 – precisely as our intuitions decree.

---

[11]   Campbell Brown considered a theory along these lines in his PEA Soup post on "Non-consequentialist utilitarianism."

Now, it is possible that, in classifying Direct Utilitarianism as a consequentialist theory, competent speakers are simply misapplying the term. In that case, our intuitions here would not count at all against D2 as a definition for consequentialism. The plausibility of this diagnosis may be enhanced by the fact that Direct Utilitarianism is weird. Isn't it possible that we are simply *confused* when we think about how to classify such a odd view? But consider (c) by itself, apart from (d). Suppose someone were simply to assert (c). This assertion, when separated from (d), no longer looks particularly odd; in fact, I suspect that something like (c) is, as often as not, used as a rough approximation of utilitarianism. And I take it that most of us are likely to say that (c), on its own, is sufficient for consequentialism. Yet we are entitled to say this only if we accept D1 for our definition of consequentialism; we are not entitled to say so if we accept D2 for our definition. This is because, if we accept D2 for our definition, then (c) is no longer sufficient for consequentialism (since (c) is consistent with (d), and the conjunction of (c) with (d) is, we have seen, inconsistent with D2).

4. Deontologize this

I have argued that D2 faces the problem that it misclassifies Direct Utilitarianism. The severity of this problem is debatable; after all, we are talking about *just one* misclassified moral theory. If this were D2's only crime, one might think its record quite good. In this section I want to expose what may be taken to be a more serious problem. To raise the problem, we need to introduce some jargon.

To **consequentialize** a given non-consequentialist moral theory N is to find (or invent) a consequentialist theory C that issues all the same verdicts, in all the same (real or hypothetical) circumstances, as N. The **Deontic Equivalence Thesis** (**DET**) is the claim that every non-consequentialist moral theory can be consequentialized. DET is true exactly if there are no uniquely non-consequentialist verdicts (or sets of verdicts), i.e. no verdicts (or sets of verdicts) unavailable to any consequentialist.[12]

The possibility and significance of DET has been widely discussed and debated.[13] Some think DET is false. Others think DET is true, and implies that all moral theories are consequentialist theories. Still others think that DET is true, and that DET is compatible with the existence of genuinely non-consequentialist theories, but that DET implies that consequentialism is necessarily true or trivial (so any genuinely non-consequentialist theory must be false or incoherent). And of course it is possible to think that DET is true, but has none of these sorts of implications. Interestingly, however, no one (to my knowledge) has noticed an alternative possibility, which I will call **Reverse Deontic Equivalence Thesis** (**RDET**).

Let us say that to **deontologize** a given consequentialist theory C is to find (or invent) a consequentialist theory N that issues all the same verdicts, in all the same (real or hypothetical) circumstances, as C. RDET, then, would be the claim that every consequentialist theory can be deontologized. RDET is true if and only if there are no uniquely consequentialist verdicts (or sets of verdicts), i.e. no verdicts (or sets of verdicts) unavailable to any non-consequentialist.

---

[12]   See Dreier (1993) for the most influential discussion of consequentialization. The term "Deontic Equivalence Thesis" is (I believe) coined in Portmore (2007).

[13]   See Dreier (1993), Brown (2004), Louise (2004), Portmore (2007), Schroeder (2007).

I want to show that RDET is an implication of D2; I also think RDET is obviously false. Thus, I argue, we must reject D2. Further, I want to show RDET is *not* an implication of D1; thus I think D1 has a major advantage over D2.

Let us assume D2 and see whether we can verify RDET. To do so, we would need to find a general procedure that we could use to deontologize any given consequentialist moral theory. The example of Direct and Indirect Utilitarianism, considered in the previous section, provides the clue to such a general procedure. Assuming D2, Direct Utilitarianism is a deontologization of Indirect Utilitarianism. The difference between them is quite simple; it is (roughly) as follows. For the Indirect Utilitarian, *pleasure* (or pain) makes consequences *good* (or bad), and *goodness* (or badness) makes actions *right* (or wrong). Goodness functions as a middleman, for the Indirect Utilitarian. The Direct Utilitarian simply skips the middleman, but delivers the same product: for the Direct Utilitarian, *pleasure* (or pain) *directly* makes actions *right* (or wrong). According to D2, goodness (or value) must always play this sort of role for any genuinely consequentialist moral theory. And this opens the possibility that, for any genuinely consequentialist moral theory, we will always be able to produce a non-consequentialist moral theory to "mimic" it, i.e. to yield the same verdicts in the same (real or hypothetical) circumstances – to *deontologize* the original consequentialist theory.

These considerations suggest that – assuming D2 – any genuinely consequentialist moral theory must be deontologizable. This is not so for D1. For, as we have seen, D1 does not require consequentialists to assign the role of middleman to value or goodness. Direct Utilitarianism is, according to D1, every bit as much a version of consequentialism as Indirect Utilitarianism. As a result, the path to deontologization is not so easy if we assume D1.

Thus, I propose that while RDET is very likely to be true given D2, it is not likely to be true given D1. This is significant because RDET is deeply counterintuitive: we normally think that there *are* uniquely consequentialist sets of moral verdicts. Consider a standard case. If a person says that a doctor is obligated to kill an innocent patient with a headache in order to harvest the patient's organs and translplant them into five sick patients who would otherwise die – and further claims that a similar obligation would obtain in any relevantly similar scenario – then we normally take ourselves to be justified in calling this person a consequentialist. But if RDET is true, then we would not be so justified. Indeed, given RDET, it would not matter how many cases one considers and makes this sort of judgment about; one will never, by that method, commit oneself to consequentialism. This is so far away from the way philosophers actually use the term "consequentialism" that I think we must conclude RDET is false – and along with it, the descriptive definition provided by D2.

5. Consequentialism and common-sense morality

In the previous sections, I have argued in favor of D1 as a definition for consequentialism. Now I will argue that, *given* D1 as a definition for consequentialism, consequentialists cannot express agreement with certain judgments that seem part-and-parcel of common-sense morality.

Consider two cases:

Susan-lie. At noon, Susan promises Joe that she will meet him for coffee at 5pm. However, when 5pm arrives, Susan simply does not bother to go to the coffee shop.

Susan-no lie. At noon, Joe invites Susan to meet him for coffee at 5pm. Susan says she may come, but warns Joe not to count on it. When 5pm arrives, Susan simply does not bother to go to the coffee shop.

According to common-sense morality, Susan's 5pm action (or, if you like, inaction) is (all else equal) wrong in *Susan-lie*, but not wrong in *Susan-no lie*. Can consequentialists agree? The only stated difference between the two cases is that Susan makes a promise at noon in *Susan-lie*, whereas Susan does not make that promise in *Susan-no lie*. But making this promise is obviously not a consequence of Susan's action: an event at noon cannot be a consequence of an action at 5pm. Indeed, for all we know, the *consequences* of Susan's 5pm action, in both scenarios, are precisely the same. Yet given D1, a consequentialist could not explain a moral difference between the two scenarios without being able to cite a corresponding difference in the consequences. Thus, it seems that (assuming D1) the consequentialist cannot justify agreement to the common-sense moralist's differing verdicts.

James Dreier (1993) has suggested an important counterargument to this view. It seems that one of the consequences of Susan's action in *Susan-lie* is that *a promise has been broken.* Yet Susan's action in *Susan-no lie* does not have this consequence. Might not the consequentialist use this difference to explain why Susan's action is wrong in one case but not the other?

To begin with, however, it is not very precise to say that one of the consequences of Susan's action is that "a promise has been broken." When this is said about the scenario in *Susan-lie*, what one really means is that (a) a promise Susan made at noon to meet Joe at 5pm is broken at 5pm. But (a) is equivalent to the following conjunction: that (b) at noon, Susan made a promise to meet Joe at 5pm, yet (c) at 5pm, Susan did not meet Joe. But it would be very odd to assert that this conjunction is a consequence of Susan's action at 5pm. *One* of the conjuncts, i.e. (c), is a consequence of Susan's action; (b) clearly isn't. Therefore, despite appearances, it is not true that one of the consequences of Susan's action (in *Susan-lie* but not in *Susan-no lie*) is that a promise is broken. Instead, the consequence (in both scenarios) is just that Susan fails to meet Joe at 5pm; the fact that this consequence is a promise-breaking in one scenario, but not the other, is not itself a consequence of Susan's action.

I suggest, therefore, that in this case, the consequentialist cannot consistently agree with the common-sense moralist – as long as D1 is assumed. However, it bears noticing, if we were entitled to switch from D2 to D1, the consequentialist's task becomes considerably easier. Susan's 5pm action (or inaction) may have the same consequences in both *Susan-lie* and *Susan-no lie*; but that does not commit us to say that the consequences *have the same value* in both cases. Joe is sitting by himself at 5:15pm; in both scenarios, this is a consequence of Susan's 5pm action. But perhaps it is a *worse* consequence in *Susan-lie*, in virtue of Susan's having promised to be there, than it is in *Susan-no lie*, in which no promise was made. If so, then (assuming D2) this *evaluative*

difference could be used to explain the *moral* difference between the two cases. Unfortunately, this option is not available to us, if my argument in the previous sections has been successful.

6. Conclusion: Consequentialism, value, and time

I have defended D1 on the ground that it is more consistent than D2 with most philosophers' patterns of usage of "consequentialism." In particular, D1 does a better job than D2 of reflecting our judgments about whether various moral theories are consequentialist. In this sense, D2 is more *revisionary* than D1. This might be taken to be one more count against D2. But it is worth noticing that D1 is, in another sense, more revisionary than D2. Since D1 omits any mention of *value, goodness,* and the like, it is, in this respect, a departure from a view of consequentialism which is very widespread: a view according to which consequentialism is, at bottom, a view about how the *deontic* relates to the *evaluative*. If my argument is correct, then this widely accepted characterization of consequentialism is basically mistaken.

If we are to cease thinking of consequentialism as a view about how the evaluative and the deontic relate to one another, how *should* we think about it? Superficially, if we accept D1, then consequentialism is a view about how the deontic relates to *consequences*. But what are consequences? This is a terribly difficult question. Perhaps consequences are, as some maintain, causal effects, or perhaps they are worlds, as others maintain. My view is that a necessary condition for E to be a consequence of A is that E be an *event* that occurs *after* A. Further, I think this condition is the *distinctive* characteristic of a consequence; it is the primary way in which one thing becomes *eligible* to be a consequence of another thing. (I will not argue these points at any length here, but the seeds of my argument for it are contained in Section 5.) If this is right, then it is, according to consequentialism, a *temporal relation* which qualifies an event E for relve o

References

Foot, Philippa (1983)
"Utilitarianism and the Virtues." *Proceedings and Addresses of the American Philosophical Association* 57(2)

Schroeder, Mark (2007)
"Teleology, Agent-Relative Value, and 'Good'" *Ethics* 117(2) 265-295

Portmore, Douglas (2007)
"Consequentializing Moral Theories." *Pacific Philosophical Quarterly* 88 39-73

Anscombe, GEM (1958)
"Modern Moral Philosophy." *Philosophy* 33.

Darwall, Stephen (2002)
Introduction to *Consequentialism*. Oxford: Blackwell.

Kagan, Shelley (1989)
The Limits of Morality. Oxford: Blackwell

Slote, Michael (1985)
Common-Sense Morality and Consequentialism London: Routledge.

Shaw, William (2006)
"The Consequentialist Perspective," in *Contemporary Debates in Moral Theory* (ed. James Dreier)

Sinnott-Armstrong, Walter (2007)
"Consequentialism," *Stanford Encyclopedia of Philosophy*.

Hare, RM (1952)
The Language of Morals.

Dreier, James (1993)
"Structures of Normative Theories," *The Monist* 76

Thomson, Judith (1994)
"Goodness and Utilitarianism," *Proceedings and Addresses of the American Philosophical Association* 67.

Brown, Campbell (2004).
"Consequentialise This." Unpublished manuscript.

Louise, Brown (2004)

"Relativity of Value and the Consequentialist Umbrella." *Philosophical Quarterly* 54.