

CONSEQUENTIALIZING MORAL THEORIES

I. Consequentializing

A. How to Consequentialize

The recipe for consequentializing a nonconsequentialist theory is simple: Take whatever considerations that the nonconsequentialist theory holds to be relevant to determining the deontic status of an action and insist that those considerations are relevant to determining the proper ranking of outcomes. In this way, the consequentialist can produce an ordering of outcomes that when combined with her criterion of rightness yields the same set of deontic verdicts that the nonconsequentialist theory yields—that is, for any deontic predicate ('permissible', 'impermissible', 'obligatory', or 'supererogatory'), the resulting consequentialist counterpart theory and the original nonconsequentialist theory will be in perfect agreement as to the set of actions that are in the extension of that predicate (Brown 2004, 2).

To illustrate how consequentializing works, suppose that a nonconsequentialist holds that an agent's intentions are relevant to determining whether or not her actions are morally permissible. More specifically, suppose that this nonconsequentialist holds that although it is, say, permissible to kill one as a merely foreseen but unintended consequence of doing what will save five others, it is impermissible to intend to kill one for the sake of saving five others. In this case, all the consequentialist needs to do to yield these same two deontic verdicts is to hold that although the outcome in which an agent has killed someone as a merely foreseen but unintended consequence of saving five others ranks higher than the outcome in which the agent has instead allowed those five others to die, the outcome in which an agent has intentionally brought about someone's death in order to save five others ranks lower than the outcome in which the agent has instead allowed those five others to die.

More abstractly, if an agent has the option of performing a_1 , a_2 , a_3 , a_4 , a_5 , or a_6 , and the non-consequentialist holds that a_1 , a_3 , and a_5 are morally permissible and that the rest are morally impermissible, then the consequentializer need hold only that o_1 , o_3 , and o_5 are tied for first place in the ranking of outcomes and that o_2 , o_4 , and o_6 each rank lower than these three top-ranked outcomes. If the non-consequentialist holds that a_4 is morally obligatory, then the consequentializer need hold only that o_4 outranks all the alternative outcomes. Of course, the ranking of outcomes will need to be agent-relative. I'll have more to say on this in a moment.

B. Why Consequentialize? The Consequentializing Project (a.k.a. the Consequentialization Program)

The project is to advance the cause of act consequentialism by demonstrating that there is a version of act consequentialism (a) that, unlike *non*-act-consequentialism, is compatible with the act utilitarianism's compelling idea and (b) that, unlike other (agent-neutral) versions of act consequentialism, is compatible with our most firmly held moral convictions, such as our conviction that there are various agent-centered options and agent-centered constraints. Such a version of act consequentialism would have significant advantages over its rivals: both its *non*-act-consequentialist rivals and its agent-neutral act consequentialist rivals. It would be a hybrid that possessed the best features of each breed.

There is something compelling about act utilitarianism even though it has wildly counterintuitive implications. The consequentializer hopes to capture what is compelling about act utilitarianism while avoiding its counterintuitive implications.

C. What's required for the consequentializing project to succeed?

For this project to succeed, the consequentializer must identify both some compelling idea and some version of act consequentialism such that both:

- (1) This version of act consequentialism is compatible with our most firmly held moral convictions, such as our conviction that there are various agent-centered options and agent-centered constraints, and
- (2) This compelling idea is something that both act utilitarianism and this version of act consequentialism are compatible with, but that no *non*-act-consequentialist theory is compatible with.

II. Two Failures

A. The First Failure: Agent-Relative Teleology and the Compelling Idea

1. Agent-Relative Teleology

ART For all agents x , x ought always to perform the act that will produce the state of affairs that, of all those available to x , is best-relative-to x .

The ART-ist can accommodate our intuition that it would be wrong for Franz to commit murder in order to prevent both Hans and Jens from each

committing murder as well as wrong for Jens to commit murder in order to prevent both Hans and Franz from each committing murder. To do so, the ART-ist need only claim that whereas the state of affairs in which Hans and Jens each commit murder is worse-relative-to Franz than the state of affairs in which Franz commits murder, the state of affairs in which Hans and Franz each commit murder is worse-relative-to Jens than the state of affairs in which Jens commits murder.

ART and act utilitarianism are structurally similar in that they both tell agents, first, to rank all the outcomes that they are capable of producing and, second, to perform an act that will produce one of the top-ranked outcomes, but whereas utilitarianism ranks states of affairs according to the *better than* relation, ART ranks states of affairs according to the *better-relative-to* relation. Unlike the *better than* relation, which is just a two-place relation between two states of affairs, the *better-relative-to* relation is a three-place relation between two states of affairs and an agent. It is the fact that this *better-relative-to* relation yields rankings that are relativized to the agent that provides ART with its ability to accommodate agent-centered constraints.

2. How are we supposed to understand the *better-relative-to* relation?

There are two possibilities. One possibility is that we have no pre-theoretical grasp of the relation (just as we have no pre-theoretical grasp of what an electron is), and so everything we know about the relation is given by the theory itself in which it functions (viz., ART). The other possibility, of course, is that we do have some pre-theoretical grasp of the relation such that we can understand the relation prior to, and independently of, having any understanding of ART.

3. Why ART-ists want (and, indeed, need) to claim that we have some pre-theoretical grasp of the *better-relative-to* relation:

To see why, consider the alternative, the alternative being that the *better-relative-to* relation is just some theoretical posit. If the *better-relative-to* relation is just some theoretical posit on the model of an *electron*, then it shouldn't matter what we call it. We might just as well have used the words *oranger-relative-to* and *orangest-relative-to* in place of *better-relative-to* and *best-relative-to*. And if it doesn't matter what we call this purely theoretical relation, then Agent-Relative Orangeology should be equally attractive as Agent-Relative Teleology.

ARO For all agents x , x ought always to perform the act that will produce the state of affairs that, of all those available to x , is orangest-relative-to x .

This will not do for the ART-ist's purposes, for the ART-ist thinks that ART is more attractive than ARO. The ART-ist doesn't want merely to accommodate agent-centered constraints, as ARO can; the ART-ist wants also to accommodate whatever it is that is supposedly compelling about act utilitarianism, which, according to many ART-ists, is the Compelling Idea:

CI For all agents x , x is always morally permitted to perform the act that will produce the state of affairs that, of all those available to x , is best.

ARO, clearly, doesn't accommodate CI, because the *oranger-relative-to* relation is just a theoretical posit that bears no resemblance to the *better than* relation. But the real question is whether ART can accommodate CI. It is not clear that it can, for it seems that ART implies, not the Compelling Idea (CI), but instead a Potentially Different Idea:

PDI For all agents x , x is always morally permissible to bring about the state of affairs that, of all those available to x , is best-relative-to x .

The hope, then, is that, given the right sort of theory about the semantics of 'good', 'better', and 'best', it will turn out that CI just means PDI and is, therefore, compatible with ART. Mark Schroeder (2007), however, has argued quite convincingly both (1) that ART-ists have not actually offered us any theory of the semantics of 'best' according to which CI just means PDI and (2) that even if they did, it couldn't be the correct theory: any creative semantics that yielded the result that CI just means PDI would have to change the subject from the ordinary word of English 'best' to some theoretical notion that just happens to use the letters b-e-s-t in its name.

4. Why ART-ists have not actually offered us any theory of the semantics of 'best' according to which CI just means PDI and is, therefore, compatible with ART:
 - a. The Theory that ART-ists Have Offered: The Indexicalist Theory

According to Sen and Portmore and Smith, the connection between 'good' and the *good-relative-to* relation is that 'good' expresses the *good-relative-to* relation in a context-dependent way. For example, according to Smith, "If this is right, however, then, as is perhaps already clear, it turns out that 'good' is indeed subscripted in just the way required. For when I judge *p*'s being the case in *C* to be good, I am judging that *p*'s being the case in *C* has a certain relational property. . . . In other words, I am really judging *p*'s being the case in *C* to be good_{me} and you are really judging *p*'s being the case in *C* to be good_{you}. Our judgments are appropriately relational." Sen and Portmore also tell us that the truth of evaluations (sentences calling something good) is relative to the evaluator (to the speaker). The concrete view that these remarks most strongly suggest is the view that 'good' is an *indexical*, picking out the relational property, *good-relative-to x*, where *x* is the speaker of the context. (Schroeder 2007, 283)

On this indexicalist theory, when Franz says, "my murders are worse than Hans's" and Hans says, "my murders are worse than Franz's," they both speak truly, for Franz is saying, "my murders are worse_{Franz} than Hans's" and Hans is saying, "my murders are worse_{Hans} than Franz's," and these two don't contradict each other.

This is the only theory that ART-ist have offered, so let's consider whether it will do the work that the ART-ist needs it to do.

- b. Why the indexicalist theory fails to establish that CI just means PDI and is, therefore, compatible with ART:

Now, one problem with this view is that it is implausible. But, as Schroeder points out, even if it were plausible, it doesn't get the ART-ist what she needs.

According to the indexicalist theory, when Doug says, "ART is compatible with CI," he is saying that the following two propositions are compatible:

- (a) For all agents *x*, *x* ought always to perform the act that will produce the state of affairs that, of all those available to *x*, is best_{*x*} (i.e., best-relative-to *x*).
- (b) For all agents *x*, *x* is always morally permitted to perform the act that will produce the state of affairs that, of all those available to *x*, is best_{Doug}.

But (a) is not compatible with (b), for (b) implies that it is always permissible for Franz to bring about what's best_{Doug}, whereas (a) implies that it is always permissible for Franz to bring about what's best_{Franz}. So when anyone says, "ART is compatible with CI," she speaks falsely. Thus the claim is false.

On the indexicalist theory, PDI means (b)—or something analogous to (b) with the speaker's name substituted for 'Doug'—and yet (a) is, clearly, not equivalent to PDI. Therefore, CI doesn't just mean PDI on the indexicalist theory. And since the indexicalist theory is the only theory that ART-ists have offered, we can conclude that ART-ists have not actually offered us any theory of the semantics of 'best' according to which CI just means PDI.

Now I'm not sure that the indexicalist theory is the most charitable way to be interpreting what the ART-ists are saying, but, as Schroeder will argue, it doesn't matter: any creative semantics of 'best' that ART-ists might give on which CI just means PDI would be one on which 'best' means something other than what ordinary speakers of English (including moral philosophers) have meant by it.

5. Why, even if they had, it couldn't be the correct theory of the semantics of 'best':

To accept some view [such as CI], you have to accept the sentence expressing that view on the semantics *on which it expresses that view*. Any atheist can accept the sentence 'God exists', if given a semantics on which 'God' has my wristwatch as its referent, or on which 'exists' is synonymous with 'is preposterous'. This does not make them theists. To be theists, they must accept the sentence 'God exists' under the semantics on which it means that God exists. (Schroeder 2007, 281)

There is compelling evidence that any creative semantics of 'best' on which CI just means PDI would be one on which 'best' means something other than what ordinary speakers of English (including moral philosophers) have meant by it. Here's the evidence, as offered by Schroeder (2007):

1. Agent-centered constraints are widely supposed by competent speakers of English to be counterexamples to CI.
2. Competent speakers of English take CI to imply that it "is always morally permissible for everyone to do what will have results that rank highest in the *better than* ordering," and they take this to imply that there "is an ordering, the *better than* ordering, such that it is always permissible for everyone to do what will have the results that rank

highest in it" (Schroeder 2007, 281). But this last claim is inconsistent with CI.

6. Why CI can't be act utilitarianism's compelling idea:

This would all be pretty devastating with regard to the prospects for the success of the consequentializer's project were it not for the fact that CI cannot be act utilitarianism's compelling idea.

It is generally recognized that act utilitarianism's compelling idea is something that is compatible with ethical egoism—see, for instance, Scheffler (1985, 414 & 416), Portmore (2007), Shaw (2006, 16), Schroeder (2007, 291, note 44). But CI is not compatible with ethical egoism. Ethical egoism denies that agents are always morally permitted to act so as to bring about the best available state of affairs, for ethical egoism entails that it would be wrong to so act when doing so would involve failing to bring about the state of affairs that would be *best for* the agent.

B. The Second Failure: Dual-Ranking Act Consequentialism and the Permissibility of Maximizing View

1. Dual-Ranking Act Consequentialism (DRAC)

Dual-Ranking Act Consequentialism: S's performing x is morally permissible if and only if there is no available act alternative that would actualize a possible world that S has both (i) more moral reason to want to be actualized than to want Wx to be actualized and (ii) at least as much reason, all things considered, to want to be actualized as to want Wx to be actualized. And S's performing x is supererogatory if and only if there exists an alternative act y such that y is morally permissible and S has more moral reason to want Wx to be actualized than to want Wy to be actualized.

Let Wx be the possible world that would be actualized were S to perform x .

One substantive version:

Self/Other Utilitarianism (SOU): S's performing x is morally permissible if and only if there is no available act alternative that would produce both (i) more utility for others (i.e., those other than S) than x would and (ii) at least as much overall utility as x would. And S's performing x is supererogatory if and only if it produces more utility for others than some other

permissible alternative would. (This is adapted from Sider, "Asymmetry and Self-Sacrifice," p. 128)

Let $U_s(x)$ = the utility that accrues to S if S does x , $U_{-s}(x)$ = the utility that accrues to others if S does x , and $U(x)$ = the overall utility that is produced if S does x .

<i>act</i>	$U_s(x)$	$U_{-s}(x)$	$U(x)$	<i>moral status</i>
a1	+70	+10	+80	merely permissible
a2	+20	+5	+25	impermissible
a3	+10	+15	+25	supererogatory
a4	-10	+20	+10	supererogatory and morally best

2. The Permissibility of Maximizing View (PMV)

PMV S is always morally permitted to perform the act that would actualize the possible world that S has most reason to want to be actualized.

3. Why the Consequentializing Project Fails Given These Identifications

PMV is an idea that a *non-act-consequentialist* can accept. To accept this, a *non-act-consequentialist* need only accept the following thesis:

FT* S has most reason to want Wx to be actual if and only if S morally permitted to perform x .

As I said above, for the consequentializing project to succeed, the consequentializer must identify both some compelling idea and some version of act consequentialism such that both:

- (1) This version of act consequentialism is compatible with our most firmly held moral convictions, such as our conviction that there are various agent-centered options and agent-centered constraints, and
- (2) This compelling idea is something that both act utilitarianism and this version of act consequentialism are compatible with, but that no *non-act-consequentialist* theory is compatible with.

Although DRAC does meet condition (1), and although both act utilitarianism and DRAC are compatible with PMV, it is not the case that no *non-act-consequentialist* theory is compatible with PMV. Those versions of *non-act-consequentialist* that endorse FT* are compatible PMV.

Thus, the consequentializing project fails if the compelling idea is PMV.

3. *Non-Act-Consequentializing*

What's even worse is, as Sachs argues, that "for any set of deontic verdicts, Δ , Δ can be delivered by a non-[act-]consequentialist theory, N, that is compatible with PMV."

In order to consequentialize N with deontic verdicts Δ , the consequentializer must put forward some agent-relative ranking of possible worlds in terms of how much reason the agent has to desire their actualizations such that Δ is compatible with PMV—call this ranking R. So, for instance, if, according to N, x is permissible and y is impermissible, then the consequentializer must put forward a ranking where Wx outranks Wy . Now, by stipulation N is compatible with Δ . So, assuming that the *is-compatible-with* relation is a transitive relation, we can conclude that N is compatible with PMV given R.

So there is a version of consequentialism, C, that is compatible with both PMV and Δ only if R. But if R, N is also compatible with PMV. And, in that case, consequentializing doesn't advance the cause of consequentialism, because there is a C that's compatible with both PMV and Δ only if there is an N that's compatible with both PMV and Δ . But, then, it seems that C has nothing over N. They both yield the same set of deontic verdicts, Δ , and they are both compatible with the compelling idea called PMV.

Of course, there may be some Ns for which this won't work, because some Ns might be committed to a view that implies $\sim R$. Assume that the rule consequentialist is committed to the view that the better a world the more reason everyone has to want it to be actualized. If this is right, then the rule consequentialist can't help herself to R to show that her theory is compatible with PMV, because the rule consequentialist must, in that case, reject R.

Nevertheless, as Sachs argues, it is still true that, "for any set of deontic verdicts, Δ , Δ can be delivered by a non-[act-]consequentialist theory, N, that is compatible with PMV." For instance, there will be an N that just consists of a list of DOs and DON'Ts, one for every member of Δ .

Sachs is right, of course, but it would seem that C would have a huge advantage over such a theory. N and C would both be compatible with Δ and PMV, but C would be superior in its ability to systemize our moral convictions, as set forth in Δ .

4. Why Sachs thinks that no matter what we identify as being act consequentialism's compelling idea and no matter what we identify as being the set of deontic verdicts that a theory should ideally comport with, there will be a *non-act-consequentialist* theory that is compatible with both:

Sachs claims: For any plausible candidate for consequentialism's compelling idea, *cI*, and for any set of deontic verdicts, Δ , if there is an act consequentialist theory, *C*, that can deliver Δ while remaining compatible with *cI*, then there is a *non-act-consequentialist* theory, *N*, that can do the same.

Sachs proposes the following method for determining whether there is any *cI* that would serve as a counterexample to the above universal claim: "we will examine one-by-one the four elements of non-consequentialist theories—the four kinds of claim that non-consequentialist theories sometimes make—to determine whether any of them would be incompatible with any likely compelling idea. Those four elements are: First, one or more moral principles. Second, evaluative claims. Moral principles, combined with non-moral facts (which themselves, I take it, are not part of normative ethical theories) and/or evaluative claims, yield the third part of a normative ethical theory: a set of deontic verdicts. The fourth part is a determination claim—a claim about what kinds of facts ground, or determine, the truth of the deontic verdicts."

5. Where Sachs Goes Wrong

Sachs admits that if *cI* is a determination claim, then this would prove to be a counterexample to the above universal claim, but he doubts that *cI* is a determination claim. In support of this, he says, "consequentialism's determination claims are usually seen as one its main drawbacks. If Jones stalks and kills an innocent person in cold blood, consequentialism tells us that it's the resulting state of affairs that determines whether Jones's action was impermissible. Thus, even if consequentialism gives the right answer, it will seem to give it for the wrong reason." I'm not convinced that consequentialism does give the right answer for the wrong reason, but, for the sake of argument, let's assume that's true. Even so, why can't act consequentialism's determination claim be both what's compelling about it and also what's troubling about it? Why can't it be compelling in that it is entailed by moral rationalism and the teleological conception of practical reasons, but troubling in that it doesn't jibe with our intuitions about what makes it wrong to stalk and kill Jones?

Also, Sachs wrongly assumes that the only thing that could be incompatible with non-act-consequentialism's determination claim is act consequentialism's determination claim. Thus the only thing he considers with regard to the issue of whether there is some cI that is incompatible with N's determination claim is C's determination claim. But if cI is some conjunctive claim (e.g., MR + TCPR) that implies that N's determination claim is false, then cI would still be a counterexample to the above universal claim. So Sachs needs to argue that no set of claims that implies C's determination claim could be cI. And he hasn't done so. This is especially problematic, because, as I've argued in an earlier lecture, it is MR+TCPR that people like Scheffler (and Shaw) have taken to be cI.

III. One Big Success

A. Success for the Consequentializing Project: Dual-Ranking Act Consequentialism (DRAC) and Act Consequentialism (AC)

DRAC S's performing x is morally permissible if and only if there is no available act alternative that would actualize a possible world that S has both (i) more moral reason to want to be actualized than to want Wx to be actualized and (ii) at least as much reason, all things considered, to want to be actualized as to want Wx to be actualized. And S's performing x is supererogatory if and only if there exists an alternative act y such that y is morally permissible and S has more moral reason to want Wx to be actualized than to want Wy to be actualized.

AC An act has the moral status that it does solely in virtue of what reasons there are for and against preferring its outcome to those of its available alternatives, such that, if S is morally required to perform x , then, of all the outcomes that S could bring about, S has most reason to want Wx to obtain.

As we saw in an earlier lecture, it is plausible to suppose that AC (or, alternatively, MR+TCPR) is act utilitarianism's compelling idea. And we explained why AC is so compelling in terms of two other compelling ideas: moral rationalism and the teleological conception of practical reasons. The argument went as follows:

1. An act has the moral status that it does solely in virtue of what reasons there are for and against performing it, such that, if S is morally required to perform x , then S has most reason to perform x . (from Strong Moral Rationalism)

2. The reasons there are for and against performing a given act are what they are solely in virtue of what reasons there are for and against preferring its outcome to those of its available alternatives, such that, if S has most reason to perform x , then, of all the outcomes that S could bring about, S has most reason to want Wx to obtain. (from the Teleological Conception of Practical Reasons)
3. Therefore, an act has the moral status that it does solely in virtue of what reasons there are for and against preferring its outcome to those of its available alternatives, such that, if S is morally required to perform x , then, of all the outcomes that S could bring about, S has most reason to want Wx to obtain. (Maximizing Act Consequentialism)

Now, AC is an idea that both act utilitarianism and DRAC are compatible with, and that no *non-act-consequentialist* theory is compatible with—clearly no *non-act-consequentialist* theory can endorse AC. What’s more, DRAC is a version of act consequentialism that is compatible with our most firmly held moral convictions, such as those deriving from agent-centered options and agent-centered constraints. Therefore, the consequentializing project has succeeded. We’ve identified both a compelling idea (*viz.*, AC) and a version of act consequentialism (*viz.*, DRAC) such that both:

- (1) This version of act consequentialism is compatible with our most firmly held moral convictions, such as our conviction that there are various agent-centered options and agent-centered constraints, and
- (2) This compelling idea is something that both act utilitarianism and this version of act consequentialism are compatible with, but that no *non-act-consequentialist* theory is compatible with.

B. How AC is compatible with both act utilitarianism and ethical egoism

Recall that it is generally recognized that act utilitarianism’s compelling idea is something that is compatible with both act utilitarianism (AU) and ethical egoism (EE). But it may seem that AC is not compatible with either AU or EE. Take AU. AU says that the moral status of an act (*i.e.*, M) is determined by the goodness of its outcome (*i.e.*, G), whereas AC says that the moral status of an act (*i.e.*, M) is determined by how much reason the agent has to desire that its outcome obtains (*i.e.*, D). But how can M be determined by D, as AC claims, if M is determined by G, as AU claims? The answer, I think, is that it is plausible to suppose that G is determined by D, for, if M is determined by G and G is determined by D, then it follows that M is determined by D. In fact, this will be the case if the buck-passing account of value is true. Now, since AU is compatible with the buck-passing account of value, and since AU is just a version of AC if the buck-passing account of ‘good’ is correct, it follows that AC

is compatible with AU. Likewise, AC is compatible with EE, since EE is compatible with the buck-passing account of 'good for', and since EE just is a version of AC if the buck-passing account of 'good for' is correct.

The thought here is that what makes the following types of claims all evaluative claims is that they are each to be understood as claims about what there are reasons to prefer.

Better that Claims: For all states of affairs p and q , it is better that p is the case than that q is the case if and only if the set of all the right kind of reasons to prefer its being the case that p to its being the case that q is weightier than the set of all the right kind of reason to prefer its being the case that q to its being the case that p .

Better for Claims: For all subjects S and all states of affairs p and q , its being the case that p is better for S than its being the case that q is if and only if the set of all the right kind of reasons to prefer its being the case that p to its being the case that q for S 's sake is weightier than the set of all the right kind of reason to prefer its being the case that q to its being the case that p for S 's sake. (Adapted from Schroeder, "Value Theory," SEP)

Better K Claims: For all kinds K and all things A and B , A is a better K than B is if and only if the set of all the right kind of reasons to prefer A to B when selecting a K is weightier than the set of all the right kind of reasons to prefer B to A when selecting a K . (Adapted from Schroeder, "Value Theory," SEP)

If this is right, then what EE and AU have in common is that they both make what an agent ought to do a function of what reasons she has to prefer certain outcomes to their available alternatives. Thus if we want to define act consequentialism such that EE and AU are both species of act consequentialism, then we need to define act consequentialism as follows:

AC An act has the moral status that it does solely in virtue of what reasons there are for and against preferring its outcome to those of its available alternatives, such that, if S is morally required to perform x , then, of all the outcomes that S could bring about, S has most reason to want Wx to obtain.

C. The Deontic Equivalence Thesis (DET) – a.k.a. Dreier's Conjecture

DET For any remotely plausible *non-act-consequentialist* theory, N , there is a consequentialist counterpart theory, C , that yields, in

every possible world, the exact same set of deontic verdicts that N does, including not only such verdicts as 'permissible' and 'impermissible', but also such verdicts as 'supererogatory'.

See my "Consequentializing Moral Theories" for a defense of DET.

D. From DET, does it follow that we are all act consequentialists?

I argue that this can't follow since analogues of DET are true of most moral theories, including contractualism, Kantianism, virtue theory, etc. Thus if DET establishes that we're all consequentialists, then these analogues establish that we are all Kantians, contractualists, and virtue theorists as well, and all at the same time, which is just absurd. Take Kantianism, for instance. The Kantian analogue of DET is: for any remotely plausible non-Kantian theory, M_{-K} , there is a Kantian counterpart theory, M_K , that yields, in every possible world, the exact same set of deontic verdicts that M_{-K} does. To illustrate, suppose that a Kantian takes the principle "an act is morally permissible if and only if it involves treating humanity as an end-in-itself and never simply as a means" to be the one and only fundamental moral principle. The recipe for Kantianizing is, then, as follows. Take whatever considerations M_{-K} regards as relevant to determining the deontic status of an action and insist that those considerations are relevant to determining whether humanity has been treated as an end-in-itself. So suppose that M_{-K} is traditional act-utilitarianism (TAU). The theorist who wants to Kantianize TAU need only insist that we treat humanity as an end-in-itself if and only if we give equal consideration to everyone's interests in maximizing aggregate utility. Similarly, we could follow the same procedure for any other M_{-K} .

If we are to assume that, from DET and DET alone, it follows that we're all consequentialists, then, by the same reasoning, it follows from this analogue that we are also all Kantians. But this is patently absurd, precisely because these are *rival* moral theories that accept different fundamental moral principles, which specify different fundamental right-making and wrong-making features and, hence, different rationales for the deontic verdicts that they yield.

E. From DET, does it follow that the distinction between act consequentialism and *non-act-consequentialism* is an empty distinction?

To conclude that consequentialism is empty from the mere fact that DET is true is to presume that if two moral theories are necessarily coextensive in their deontic verdicts, then there is nothing substantive at issue between them. But this is false, because moral theories do much more than just yield deontic verdicts. Importantly, they provide different competing rationales for the

deontic verdicts that they yield. Thus, as we have just seen, an act-utilitarian, a Kantian, and a contractualist can all agree that the extension of permissible acts is just those that maximize utility, but even so they will provide different explanations for why this is so, for they necessarily accept different views about what the fundamental right-making and wrong-making features of acts are. And this means that the theories, although extensionally equivalent, will have different truth conditions. For instance, AU (act utilitarianism) is true only if utility is the only thing that is good for its own sake. And utilitarian contractualism is true only if the only principle that “no one could reasonably reject as a basis for informed, unforced general agreement” is the principle of utility. And utilitarian Kantianism is true only if giving everyone equal consideration in maximizing aggregate utility is what constitutes treating humanity as an end-in-itself. These are substantive issues over which the traditional act-utilitarian, the utilitarian contractualist, and the utilitarian Kantian may disagree.

- F. An Objection to Consequentializing: Is it legitimate for the consequentialist to appeal to our moral convictions in specifying what we have reason to desire?

Since one of the consequentializer’s main aims is to avoid the counterintuitive implications that plague AU, the consequentializer will need to appeal to our considered moral convictions when determining how outcomes are to be ranked. But we might wonder whether this is legitimate. That is, we might wonder: Can a consequentialist defend one ranking over another on the basis that the former but not the latter yields intuitive moral verdicts when combined with BC?

BC An act is permissible if and only if no other available outcome ranks higher than its outcome in terms of how much reason there is for the agent to desire that the given outcome obtains.

Holding BC constant,¹ there are three possible procedures one might employ in figuring out the proper ranking of outcomes (i.e., figuring out how an agent ought to rank the various alternative outcomes that she’s capable of producing):

The Footian Procedure: Come to some fixed set of pre-theoretical judgments about the rightness of actions that is entirely independent of any judgments one has about the proper ranking of outcomes, and then hold

¹ Why hold BC constant? The idea, here, is that the act consequentialist is committed to BC. So we have to hold BC constant when determining what the most plausible substantive version of act consequentialism is. Of course, once we determine what’s the most plausible version of act consequentialism we will have to compare it to the most plausible versions of other rival moral theories and use wide reflective equilibrium to determine whether or not we should accept it.

that one outcome outranks another if and only if its corresponding act is deontically superior to that of the other, where a supererogatory act is deontically superior to any non-supererogatory alternative and a permissible act is deontically superior to any impermissible act.

The Foundationalist Procedure: Come to some fixed set of pre-theoretical judgments about the proper ranking of outcomes that is completely independent of any judgments one has about the rightness of actions, and then rank outcomes accordingly.

The Coherentist Procedure: Keeping BC constant, revise one's pre-theoretical judgments both about the proper ranking of outcomes and about the rightness of actions in light of each other until reflective equilibrium is reached, and then rank outcomes accordingly.

The Footian Procedure would not yield any normative advice beyond what our pre-theoretical deontic judgments already tell us. Thus the consequentialist cannot adopt the Footian Procedure if she wants the resulting substantive theory (the one obtained by applying the results of the Footian Procedure to BC) to yield informative results. Failing to provide any normative advice beyond what our pre-theoretical deontic judgments already tell us would be a significant failure, as one of the things that we hope to gain from moral theorizing is a deeper understanding of morality, so that we might better deal with moral questions about which we disagree or about which we have no confident pre-theoretical judgment.

But just because the consequentializer should eschew the Footian Procedure doesn't mean that she must eschew any appeal to our pre-theoretical deontic judgments and adopt the Foundationalist Procedure. Fortunately for the consequentializer, the Coherentist Procedure is a viable alternative that, like the Foundationalist Procedure, yields informative results. Moreover, the Coherentist Procedure has the advantage of being likely to yield a substantive consequentialist theory (one that combines consequentialism with a substantive account of how outcomes are to be ordered) that avoids many of the counter-intuitive implications that plague TAU.

Unlike both the adherents of the Foundationalist and Footian Procedures, the adherent of the Coherentist Procedure doesn't consider either her pre-theoretical judgments about the proper ranking of outcomes or her pre-theoretical judgments about the rightness of actions to be fixed starting points. Whereas the adherent of the Foundational Procedure will be unwilling to revise her judgments about the proper ranking of outcomes in light of any potentially counter-intuitive implications it yields when combined with BC, the adherent

of the Coherentist Procedure is willing to do so. And whereas the adherent of the Footian Procedure will be unwilling to revise her pre-theoretical judgments about the rightness of actions in light of any potential conflicts that may arise when she combines her pre-theoretical judgments about the proper ranking of outcomes with BC, the adherent of the Coherentist Procedure is willing to do so. Sometimes the adherent of the Coherentist Procedure will end up revising her deontic judgments in light of her more firmly held evaluative judgments. Other times she will revise her evaluative judgments in light of her more firmly held deontic judgments. The Coherentist Procedure will, therefore, be informative, yielding both new judgments about the deontic statuses of actions and new judgments about the how outcomes should be ranked. Given the viability of the Coherentist Procedure, we can conclude, then, that the consequentializer can appeal to our considered moral convictions when determining how outcomes are to be ordered without rendering her view either circular or uninformative.