

TRADITIONAL ACT CONSEQUENTIALISM AND THE TOO DEMANDING OBJECTION

I. What is the too demanding objection to traditional act consequentialism?

Recall that, according to *traditional act consequentialism* (TAC), an act is morally permissible if and only if it maximizes the good, and if an act is morally permissible, then what makes it so is ultimately the fact that it maximizes the good.

TAC implies that agents are morally required to sacrifice their own projects, interests, and special relationships whenever they can do more to promote the good by doing so. Thus, I am morally required to abandon my own project so as to ensure that someone else succeeds in completing hers if I could thereby produce slightly more good than I would if I didn't abandon my project. And I am morally required to sacrifice my own life to save another's if I could thereby produce slightly more good than I would if I didn't sacrifice my life. Also, I am morally required to sacrifice my relationship with my loved one to save someone else's relationship with her loved one if I could thereby produce slightly more good than I would if I didn't sacrifice my relationship with my loved one.

Now, given such implications, many object that TAC is just too demanding to be a plausible moral theory. But what exactly is the objection? The objection can't be just that TAC implies that agents are, in certain circumstances, morally required to make sacrifices, or even that they are, in certain circumstances, morally required to make tremendous sacrifices. All moral theories, except perhaps ethical egoism, entail that agents are, in certain circumstances, morally required to make tremendous sacrifices. Indeed, all plausible moral theories imply that you are, in certain circumstances, morally required to make the ultimate sacrifice: viz., to sacrifice your life. To illustrate, consider a relatively undemanding moral theory, a moral theory that includes no positive duties except those that one voluntarily takes on in virtue of entering into certain contractual relationships. Even on this sort of moral theory, a person could be morally required to sacrifice her life, as where she has entered into some contractual relationship that requires her to do so—where she has, for instance, signed on to be the President's bodyguard.

A. The Old Version of the Too Demanding Objection

Let a_1 be the act of sacrificing one's own life or the life of one's loved one (which, we'll assume, are both well worth living) for the sake of saving some stranger's life, where doing so would maximize the good, but where the net gain in the amount of good would be infinitesimally small. The old version of the too demanding objection runs as follows:

1. If TAC is true, then S is morally required to perform a₁. (From the definition of 'TAC'.)
2. S is not morally required to perform a₁. (From intuition.)
3. Therefore, TAC is false. (From 1 and 2.)

B. Problems with the Old Version

1. The old version relies on an appeal to our commonsense moral intuition concerning a particular case—see premise 2. As we've already seen, though, the consequentialist will not be impressed by such an appeal to moral intuition. The consequentialist will argue that such intuitions are not to be trusted. Of course, this doesn't mean that the argument isn't sound; it just means that it won't have much rhetorical force against the consequentialist.
2. There is, as David Sobel has argued, an even more significant problem for the old version of the too demanding objection: it relies on a nonconsequentialist premise—specifically, the premise that it's not just the size of the cost that a moral theory imposes that matters, but whether the moral theory imposes those costs by requiring agents to bear them for the sake of aiding others or imposes those costs by requiring unaided patients to bear them for the sake of others' freedom from being required to come to their aid. But if the objection presupposes some prior break from consequentialism in taking this sort of distinction to be morally relevant, then it cannot provide any independent grounds for rejecting consequentialism. Here's Sobel:

To start to make my case against the self-standing nature of the Objection, consider a different sort of situation in which people tend to feel the pull of the Demandingness Objection. Consider the case of Joe and Sally. Joe has two healthy kidneys and can live a decent but seriously reduced life with only one. Sally needs one of Joe's kidneys to live. Even though the transfer would result in a situation that is better overall, the Demandingness Objection's thought is that it is asking so much of Joe to give up a kidney that he is morally permitted to not give. The size of the cost to Joe makes the purported moral demand that Joe give the kidney unreasonable, or at least, clearly not genuinely morally obligatory on Joe. Consequentialism, our intuitions tell us, is too demanding on Joe when it requires that he sacrifice a kidney to Sally.

But consider things from Sally's point of view. Suppose she were to complain about the size of the cost that a non-Consequentialist moral theory permits to befall her. Suppose she were to say that such a moral theory, in permitting others to allow her to die when they could aid her, is excessively demanding on her. Clearly Sally has not yet fully understood how philosophers typically intend the Demandingness Objection. What has she failed to get about the Objection? Why is Consequentialism too demanding on the person who would suffer significant costs if she was to aid others as Consequentialism requires, but non-consequentialist morality is not similarly too demanding on Sally, the person who would suffer significant costs if she were not aided as the alternative to Consequentialism permits? What must the

Objection's understanding of the demands of a moral theory be such that that would make sense? There is an obvious and occupied answer—that the costs of what a moral theory requires are more demanding than the costs of what a moral theory permits to befall the unaided, size of cost held constant. The moral significance of the distinction between costs a moral theory requires and costs it permits must already be in place before the Objection gets a grip. But this is for the decisive break with Consequentialism to have already happened before we feel the pull of the Demandingness intuitions. (Sobel 2007, 3)

So, according to Sobel, the Demandingness Objection (DO) amounts to the claim that, although Argument A below is sound, Argument B below is unsound:

Argument A:

- P1) If consequentialism is true, then many typical agents are required to incur great costs even when they haven't consented to such costs (e.g., even when they haven't agreed to incur such costs).
- P2) It is not the case that many typical agents are required to incur great costs even when they haven't consented to such costs.
- C) Therefore, consequentialism is false.

Argument B:

- P1*) If nonconsequentialism is true, then many typical patients are required to incur great costs even when they haven't consented to such costs (e.g., even when they haven't agreed to incur such costs).
- P2*) It is not the case that many typical patients are required to incur great costs even when they haven't consented to such costs.
- C*) Therefore, nonconsequentialism is false.

So DO is essentially this:

- P1') If P2, but not P2*, is true, then consequentialism, but not nonconsequentialism, is too demanding.
- P2') P2, but not P2*, is true.
- C') Therefore, consequentialism, but not nonconsequentialism, is too demanding.

Sobel would claim that DO rests on P2' and that this premise rests on a prior break with consequentialism insofar as P2' presupposes that there is a morally

relevant distinction to be made between costs incurred by agents and costs incurred by patients, which is something that consequentialism denies.

C. Two New Versions of the Too Demanding Objection

Again, let a_1 be the act of sacrificing one's own life or the life of one's loved one (which, we'll assume, are both well worth living) for the sake of saving some stranger's life, where doing so would maximize the good, but where the net gain in the amount of good would be infinitesimally small. The two new versions of the too demanding objection run as follows:

1. Version 1:

1. If TAC is true, then S is morally required to perform a_1 . (From the definition of 'TAC'.)
2. S has decisive reason to refrain from performing a_1 . (From intuition.)
3. If S has decisive reason to refrain from performing a_1 , then S does not have sufficient reason to perform a_1 . (From the definitions of 'decisive reason' and 'sufficient reason'—see below.)
4. If S is morally required to perform x , then S has sufficient reason to perform x . (From moderate moral rationalism—see below.)
5. Therefore, TAC is false. (From 1-4.)

S has *decisive reason* to perform x =_{def} S is (objectively) rationally required to perform x . S has *sufficient reason* to perform x =_{def} S is (objectively) rationally permitted to perform x .

To see the force of premise 2, assume that S has a very special relationship with his daughter (one of having nurtured her from when she was only a newborn baby, one of having gone through a lot together over the years as part of the same family, one of mutual love, care, support, respect, and friendship). Assume also that S cares more about his daughter's welfare than about the stranger's welfare. Assume that S cares more about his daughter's welfare than about producing small net gains in the amount of impersonal good that there is in the world. Assume that saving his daughter is what would best satisfy his desires and is also what would best promote his self-interest. On what plausible theory of objective rationality, then, would it be rationally permissible for S to save the stranger's life instead of his daughter's? I can't think of any. Thus, in requiring S to sacrifice his daughter's life to save some stranger who would benefit only slightly more, TAC is too demanding, requiring S to make sacrifices that S has decisive reason not to make.

The thought that underlies this argument is that the fact that some system of norms yields exacting standards is objectionable only on the following two assumptions: (1) that this system of norms is such that agents always have sufficient reason to act in accordance with its standards and (2) that agents do not have sufficient reason to act in accordance with its most exacting standards. Note, for instance, that we would not object to some account of what the standards of etiquette are just because there are instances in which agents have decisive reason to violate those standards, for in the case of etiquette, we would deny (1). To illustrate, consider Dorsey's case of the diabetic:

Consider, for instance, a system of etiquette in which it is rude to begin eating before the host has served all seated guests. Assume that there is a large dinner party with a particularly slow host. Assume also that one of the first-served guests is a type-1 diabetic, who is in danger of severe hypoglycemia. Waiting until the host has served all guests would entail this person's blood sugar dropping to dangerous, hospitalizing levels. Surely, our reaction to this case is that this person should eat! But this is compatible with claiming that, in so doing, this person would be violating a norm of etiquette. In particular, one might have two possible accounts of the justifiability of the diabetic's behavior:

- A. The diabetic's behavior is rude, but she should eat because politeness, in this case, is not rationally required.
- B. The diabetic's behavior is not rude, because the relevant rule permits of a variety of exceptions, of which this case is an example.

What should we say about these possible justifications? My own view is that the A explanation is the better one. (Dorsey manuscript, 6)

Or consider the legal requirement not to speed, and suppose that you're on the way to the hospital with a loved one who's in need of immediate medical attention such that every second counts. Again, we wouldn't reject an account of the legal standards according to which there was such a requirement not to speed even under such exacting circumstances, for, in the case of legal standards, we reject (1). So the too demanding objection against a particular system of norms seems to take hold only if we assume that the system is one in which there is always sufficient reason to abide by those norms.

When it comes to morality, though, we do seem to accept (1). But we would, in the case of many moral theories, reject (2). For instance, in the case where one has signed on to be the President's bodyguard, it does seem that one has sufficient reason to sacrifice one's life to protect the President's life, as many moral theories would require. So the too demanding objection

only really takes hold against theories like TAC, where we accept both (1) and (2).

2. Version 2:

Those who think both that there is only sufficient reason, but not decisive reason, to refrain from performing a_1 and that strong moral rationalism is true will find the following version of the argument more plausible:

1. If TAC is true, then S is morally required to perform a_1 . (From the definition of 'TAC'.)
- 2*. S has sufficient reason to refrain from performing a_1 . (From intuition.)
3. If S has sufficient reason to refrain from performing a_1 , then S does not have decisive reason to perform a_1 . (From the definitions of 'decisive reason' and 'sufficient reason'.)
- 4*. If S is morally required to perform x , then S has decisive reason to perform x . (From strong moral rationalism.)
5. Therefore, TAC is false. (From 1-4.)

2* is weaker than 2, but 4* is stronger than 4. I happen to think that the two stronger claims, 2 and 4*, are true, but dualists about practical reason will deny both 2 and 4*.

D. The Contrary to Reason Objection

The new too demanding objection is really just a species of a more general sort of objection: what I'll call *the contrary-to-reason objection*. (In what follows, I'll be working with version 1 of the new too demanding objection, but the same points could be made with version 2, *mutatis mutandis*.)

Let M be some moral theory, and let a_1 be some act that M implies is morally required even though, intuitively, the agent has decisive reason to refrain from performing it.

1. If M is true, then S is morally required to perform a_1 . (From the definition of 'M'.)
2. S has decisive reason to refrain from performing a_1 . (From intuition.)
3. If S has decisive reason to refrain from performing a_1 , then S does not have sufficient reason to perform a_1 . (From the definitions of 'decisive reason' and 'sufficient reason'.)
4. If S is morally required to perform x , then S has sufficient reason to perform x . (From moderate moral rationalism.)
5. Therefore, M is false. (From 1-4.)

Note that this objection is not identical to the too demanding objection, for this objection applies not only to moral theories that require *more* self-sacrifice from us than we have sufficient reason to give, but also to moral theories that require *less* self-sacrifice from us than we have decisive reason to give. To illustrate the latter, consider a moral theory where there is an absolute moral prohibition against suicide. And consider that it is plausible to suppose that one could have decisive reason to commit suicide even when it might be less costly to oneself to refrain from committing suicide. Suppose, for instance, that the only way for me to save the billion people who I have imperiled through my negligent actions is to commit suicide. In such a case, it's plausible to suppose that I have decisive reason to commit suicide even if I would be better off, self-interestedly speaking, not doing so. If this is right, then we can object to this moral theory on the same grounds that we objected to TAC: it does, in certain circumstances, require us to act in a way that's contrary to the way in which we have decisive reason to act. To see this, let a_1 be the act of refraining from committing suicide in the case where committing suicide is the only way to save the billion people who I have imperiled through my negligent actions.

II. Moral Rationalism

The two new versions of the too demanding objection as well as the contrary-to-reason objection all presuppose some version of moral rationalism. So let's consider the various versions and whether or not they are plausible.

A. Three Versions of Moral Rationalism

1. *Weak Moral Rationalism* (WMR): If S is morally required to perform x , then S has some reason to perform x .
2. *Moderate Moral Rationalism* (MMR): If S is morally required to perform x , then S has sufficient reason to perform x .
3. *Strong Moral Rationalism* (SMR): If S is morally required to perform x , then S has decisive reason to perform x .

Note that these are not bi-conditionals. Take SMR, for instance. "Reasons of various kinds—prudential reasons, for instance—can make a particular course of action *rationaly* compelling without grounding a *moral* requirement" (Stroud 1998, 174). That is, prudential reasons can make it such that S has decisive reason to perform x without making it such that S is morally required to perform x .

Dualists about practical reason accept MMR, but reject SMR.

B. Stroud's Overridingness Thesis (OT)

SMR is close to Stroud's OT.

OT If S is morally required to ϕ , then S has most reason to ϕ .

SMR is a bit stronger than OT. SMR, but not OT, entails that one is always rationally required to act as one is morally required to act. By contrast, OT entails only that one always has most reason to act as one is morally required to act. But if we accept a satisficing version of objective rationality, one will not always be rationally required to act as one has most reason to act.

On the plausible assumption, that there is having most reason to ϕ is a necessary condition for having decisive reason to ϕ , SMR implies OT.

C. Why Stroud is confident that there is such a notion as what one has most (or decisive or sufficient) reason to do, all things considered:

"Is there really a perspective which adjudicates between the reasons generated by all the various modes of evaluation which could be brought to bear on practical predicaments? Is it really possible to evaluate one's reasons as a whole, to weigh against each other the different kinds of reasons one has? Are reasons of different types generally *commensurable*?" (175) Dualists about practical reason are dualists for precisely this reason: they deny that there is such a perspective.

The skepticism of dualists aside, one reason to say "yes" to these questions is the fact that this is precisely what we seem to do all the time. So if we do it, it must be possible to do it. And we do seem to make, on a regular basis, judgments as to what we have most reason to do, all things considered. This, indeed, seems to be the aim of our practical reasoning.

"Insofar as that seems to be what we are up to in our practical reasoning [i.e., arriving at a judgment as to what we have most reason to do, all things considered], we are committed to thinking that it makes sense to talk of what someone has most reason to do" (175).

We should also wonder why the dualist's skepticism doesn't take us down the slippery slope to absurdity. It seems that we can always make finer and finer-grained distinction between types of reasons, and there will always be some worry about whether the various distinct types of reasons are

commensurable—whether there is, for instance, some perspective that adjudicates between different kinds of moral reasons or different kinds of prudential reasons. Of course, we could just say that these perspectives are the perspectives of morality and prudence, respectively. But, then, why can't we say that it is the perspective of rationality (or of reasons-as-such) that adjudicates between all the various types of reasons?

D. Why Accept SMR (or OT)

Note that although the contrary-to-reason objection relies on only the weaker MMR, I will defend the stronger SMR. I will do so because I think that SMR is true. And note that SMR entails MMR. So, in defending SMR, I will thereby be defending MMR. Here are some reasons that Stroud cites for accepting SMR (or OT):

1. We think of morality as a constraint on the pursuit of our aims. It could not serve this felt role if it were sometimes rationally permissible (permissible, all things considered) to just ignore our moral duties when pursuing our aims.
2. We think that morality has the kind of rational authority that OT claims that it does. That is, we take moral verdicts of the form "S is morally required to ϕ " to be categorically binding and rationally overriding in a way that other practical requirements (e.g., the requirements of law, prudence, and etiquette) are not.
3. OT "reflects common convictions about sufficient reasons. We are inclined to take the fact that S was morally required to ϕ to be sufficient overall justification for her ϕ ing, even if ϕ ing was prudentially wrong, aesthetically unattractive, and so on" (177). [This actually supports only MMR, not SMR.]
4. OT is taken for granted in much of moral philosophy.

This is not a proof that OT is true, but rather only evidence that OT is part of our common conception of morality.

E. Is OT a substantive thesis or a conceptual truth?

Stroud believes that OT is a substantive thesis deserving of scrutiny. That is, Stroud believes that whether OT is true will depend on what the correct substantive accounts of rationality and morality are. Stroud offers two reasons for thinking that OT is a substantive thesis:

1. "The very fact that overridingness has come under attack in the literature in recent years shows that it is a falsifiable thesis" (178). But this falsely

presupposes that our concepts are transparent to us. Our concepts are not always transparent even to competent speakers. Consider how difficult it is, for instance, to come up with an analysis of 'God,' 'death,' 'free will,' 'knowledge', etc.

2. Secondly, Stroud says, "If – as OT seems to imply – moral requirements always generate (or at least always coincide with) the most weighty of reasons for action, that is a rather remarkable, or at least interesting, fact about our practical reasoning and about morality. At a minimum, then, there ought to be a story about *why* overridingness holds, if indeed it does" (178). But why can't that story be a conceptual one?

Curiously, though, she considers the fact that a given substantive moral theory is incompatible with OT given certain plausible assumptions about the substance of rationality to be a count against that moral theory rather than a count against OT. Although she claims that all she is trying to do is reach reflective equilibrium with regard to three interrelated issues: the correct account of moral requirements, the correct account of our reasons for actions, and the question of whether OT is true, she takes the fact that, for instance, C is incompatible with OT as a kind of *reductio* of C (187). But if OT is truly a falsifiable thesis, then the fact that an independently plausible moral theory like C is incompatible with OT is no more of a *reductio* of C than a *reductio* of OT.

F. Shafer-Landau's Presumptive Argument in Favor of Moral Rationalism

Shafer-Landau offers the following argument in favor of WMR, but I think that it can be reformulated as an argument for MMR. Here's what he says in favor of WMR:

When we deem someone's behaviour morally unjustified, we imply that he has violated a standard of appropriate conduct. Suppose such standards did not by themselves supply reasons for action. Then we would be forced to allow that though some actions are unjustified, immoral, improper, illegitimate, or inappropriate, there nevertheless may be no reason at all to avoid them. But this seems wrong—not only conceptually confused, but also gravely unfair. It seems a conceptual error to cite a standard as a guide to conduct and the basis for evaluation—say, for instance, that S ought to have ϕ -ed, and was wrong for having failed to—and yet claim that there was no reason at all for S to have ϕ -ed. (Shafer-Landau 2003, 192-193)

But I think that this line of reasoning has just as much intuitive force even if it's reformulated as an argument for MMR. It would go as follows. If MMR is false, then the fact that morality requires you to perform a certain act would not entail that this is an act that you have sufficient reason to perform. You might, then, have decisive reason to act contrary to what you're morally required to do. And, if you did, we could not legitimately fault or blame you, *qua* rational agent, for violating that moral requirement, assuming, that is, that you did so

precisely because you were aware of, and responding to, the decisive reasons that you had for doing so. It would be like faulting you for doing something rude when we knew that that was what you had decisive reason to do.

In general, it seems inappropriate to fault, blame, or criticize a person for violating some standard of conduct when that is what she had decisive reason to do and she was acting in response to her awareness of that fact. The appropriateness of criticizing someone for her actions seems to rest with her awareness of, and responsiveness to, practical reasons. If she was aware of what she had decisive reason to do and acted accordingly, we have to admit that we would have done what she did had we been in her circumstances and been both fully informed and fully rational. We can, of course, still label her action as, say, "rude," "immoral," or "illegal," but we cannot legitimately fault her as an agent. As a rational agent, she acted as she should have. So either we accept moral rationalism or we accept that we cannot always legitimately fault morally responsible agents for acting immorally.¹

Of course, no one thinks that it is always appropriate to fault an agent for acting immorally, for an agent can lack moral responsibility for her actions or for the states (e.g., the beliefs and motives) that necessitated her actions.² For instance, it would be inappropriate to criticize an agent for acting immorally if she acted immorally only because she was non-culpably ignorant of some crucially relevant fact. But if an agent is morally responsible for acting immorally, then it seems that wholehearted criticism is appropriate. Yet if moral rationalism is false, we must admit that wholehearted criticism of the immoral conduct of even those morally responsible for their conduct can be inappropriate. This suggests that there are significant costs to rejecting MMR, and thus that there is a strong presumption in favor of MMR.

G. The Importance of Not Confusing Moral Rationalism with the Thesis that Moral Reasons Are (Morally or Rationally) Overriding

Take:

- Moral Rationalism: If S is morally required to do x, then S has most reason, all things considered, to do x.

It's important not to confuse this with either of the following two theses:

¹ These thoughts are inspired by similar remarks by Shafer-Landau (2003, 192-193). But Shafer-Landau formulates moral rationalism a bit differently. On his formulation, a moral requirement to do x entails only that there is a good (but not necessarily decisive) reason to do x. Nevertheless, it seems that his presumptive argument for what he calls 'moral rationalism' counts equally in favor of my stronger MRT.

² This is meant to allow for the sorts of cases of blameless wrongdoing that Parfit discusses in *Reasons and Persons*...

- Moral Reasons Are Morally Overriding: If S has most moral reason to do x, then S is morally required to do x.
- Moral Reasons Are Rationally Overriding: If S has most moral reason to do x, then S is rationally required to do x.

Dorsey did this in an earlier version of his paper. He wrote:

“In order for morality to maintain rational supremacy, no loss of prudential value can possibly shake morality's stranglehold on rationality. The smallest moral reason must be enough to outweigh the largest possible prudential reason” (12).

This is indeed implausible. But one doesn't have to hold that moral reasons are morally or rationally overriding if one wants to accept SMR.

Indeed, I think that we should accept moral rationalism but deny that moral reasons are morally overriding. If we do, we will must hold that what morality can require of us will be limited to those acts that we have decisive reason to perform, all things considered. And this seems to be the case—at least, on commonsense morality. Consider my office hours case from my “Position-Relative Consequentialism, Agent-Centered Options, and Supererogation.”

III. Moral Rationalism and Moral Theory:

As I said, if we accept moral rationalism while denying that moral reasons are morally overriding, we will think that what morality can require of us will be limited to those acts that we have decisive reason to perform, all things considered. And if this is so, then those moral theories that require us to act in ways that we don't have decisive reason to act are mistaken. So we can use moral rationalism along with some background assumptions about rationality to test whether or not various moral theories are plausible. To see how this works, let's consider some of the moral theories that Stroud looks at:

A. Theory O:

According to O (“O” being for “other-regarding”), we ought to do what is best for others. Thus, according to O, the only thing that matters in determining the deontic status of an act is its effects on others. How the act might affect the agent's aims, interests, projects, etc. is completely irrelevant. O ignores the cost to the agent in benefiting others. If the choice is between saving my own life and saving someone from a paper cut, I ought to choose the latter.

O seems to be incompatible with OT. “It seems clear that the moral requirements set forth by O will not always have the balance of reasons on their side.... The reason for the failure of overridingness in this case is apparent. There is an entire class of reasons for action which O has left out of its calculations, namely those stemming from the *agent’s* interests, concerns, and so on. (I am assuming that any respectable theory of reasons for action has the consequence that I have at least some reason to pursue things that are in my interests, satisfy my desires, advance my aims, or something in that neighborhood...)” (180).

What we’ve learned from considering O: To the extent that a moral theory ignores an entire class of reasons for action, it will fail to be overriding—that is, it will fail to provide decisive reasons to act in accordance with its dictates.

B. Theory P:

P (the “P” being for “percentage”) does, unlike O, take the agent’s interests into account in determining an act’s deontic status. However, the agent’s interests count for only half (50%) of what anyone else’s interests count.

P also seems to be incompatible with OT, not because it ignores a whole class of reasons (*viz.*, prudential reasons), but because it fails to allot prudential reasons their actual force and weight. Surely each of us has at least as much reason to promote our own interests as to promote someone else’s, other things being equal. Yet, according to P, each of us only has half as much reason to do so.

What we’ve learned from considering P: “Overridingness will tend to fail, then, not just when a certain class of reasons has been left out of the moral calculus entirely, as with O, but also when such a calculus does not reflect the full strength of certain reasons, as with P” (181).

C. Theory C:

C (the “C” being for “consequentialism”) directs us to bring about the best consequences, weighing everyone’s interests impartially. (Stroud should really be referring to “U” for “utilitarianism,” since consequentialism doesn’t necessarily direct us to even consider people’s interests, let alone impartially.)

Stroud suggests that C’s moral requirements will fail to be overriding and for the same reason that P’s moral requirements failed to be overriding: it fails to reflect the full strength of certain reasons. My interests have special rational weight for me as compared with the interests of others. A person’s reasons to promote her own interests are generally stronger than her reasons for

promoting the interests of others (particularly, others who do not have any special relationship with her).

What we've learned from considering C: "Overridingness seems therefore to require a moral conception more sensitive to the agent's personal point of view than C" (184). This basically amounts to the negative argument against C.

The consequentialist must, therefore, either reject OT (and Stroud believes this is a high price to pay) or mount a full defense of Neutralism about reasons for action, the view that the normative strength (even if not the motivational strength) of one's reason to promote one's own interests is no greater, and in fact equal, to that of one's reason to promote the interests of any other.

D. Stroud's Conclusion:

These two alternatives for the consequentialist do violence to our ordinary conception of morality, on the one hand, and to our ordinary conception of rationality, on the other. Thus pending a refutation of OT or a defense of Neutralism, we have a strong presumption against consequentialism (i.e., utilitarianism).

IV. Can the consequentialist just deny moral rationalism?

The consequentialist could respond to these sorts of arguments by rejecting moral rationalism (this is, in fact, what Sobel does), but, as Hurley argues, this would be a Pyrrhic victory. Here's the argument:

Let E = some very exacting moral requirement that TAC entails, such as the requirement to perform a_1 , the act of sacrificing one's own life or the life of one's loved one (which, we'll assume, are both well worth living) for the sake of saving some stranger's life, where doing so would maximize the good, but where the net gain in the amount of good would be infinitesimally small.

1. If TAC is true, then E.
2. We don't have sufficient reason to act as E requires us to act.
3. Therefore, either TAC is false or there are some moral requirements (e.g., those that are very exacting, such as E) that rational agents ought to violate.

If we have to embrace the second disjunct in 3, then we must accept that we don't always have sufficient reason to act as morality requires. And if we don't always have sufficient reason to act as morality requires, then morality's role in practical reasoning will be quite minimal. In that case, we might know that we're morally required to act in some way, but that won't tell us whether we ought, all things

considered, to act in that way. It won't even tell us whether we would be rationally permitted to act in that way.

Thus saving TAC by rejecting moral rationalism (i.e., MMR) comes at a price: we'll have to accept that its exacting requirements (e.g., E) are always overruled by the requirements of practical reason. As Dorsey puts it, "To the extent that we believe that any particular obligation is too demanding, we believe that we should not have to follow that obligation—we believe that that obligation does not form part of a true account of 'how we should live'" (11).

Why is this a Pyrrhic victory? Answer: It's a Pyrrhic victory in that many consequentialists had set out to convince us not just that we are morally obligated to perform some extremely demanding actions, but that we ought (or, at least, that we have sufficient reason) to perform certain extremely demanding actions. According to Hurley, if we accept the anti-rationalist approach of defending TAC against the too demanding objection (whereby we reject MMR), then we save TAC only to rob TAC of its practical importance. As Hurley puts it,

[An anti-rationalist approach], however, only wins this strategic battle by losing the war so many consequentialists have taken themselves to be fighting. A central aim of many consequentialists has been to demonstrate that we should be doing more than moderate morality requires of us. But the approach in question lowers the bar of rational demands to act morally even as it raises the bar of moral standards.... If we accept that morality, properly understood, provides merely one among other sets of standards and that these standards lack the distinctive relationship that has been claimed for them to our reasons for acting, then morality is shifted toward the margins of meaningful inquiry into what we have good reasons to do. This would be a Pyrrhic victory, vindicating consequentialist morality only by marginalizing the role of morality in practical reason and deliberation. (2006, 705)

Imagine that after Peter Singer had finished arguing so passionately for the conclusion that we're morally obligated to give a substantial portion of our income to famine relief in his "Famine, Affluence, and Morality," he had then said: "But, of course, I'm only claiming that you *morally* ought to give that much; I admit, though, that, all things considered, you shouldn't give that much."

As a matter of fact, Singer did say something nearly as bad as this. In recounting his discussion with Thomas Nagel about this article, Singer writes:

eventually it emerged that he was assuming that if morality did demand that we give so much to famine relief, then there must be overriding reason to do so. I was making no such assumption. On my view, I could recognize that if I were totally committed to doing what I ought to do, I would give away my wealth up to the point indicated in my article; but at the same time I may, without any irrationality, choose to be less than totally committed to doing what I [morally] ought to do. My own interests, or those of my family, may counteract the demands of morality to some degree, and I may think it reasonable to give in

to them, while recognizing that it would be morally wrong for me to do so. (Singer 1999, 308-9).

It turns out that Singer denies SMR, but it seems from other things that he has said that, like Sidgwick, he accepts MMR. But this seems nearly as bad as the above. So imagine that after Peter Singer had finished arguing so passionately for the conclusion that we're morally obligated to give a substantial portion of our income to famine relief in his "Famine, Affluence, and Morality," he had then said: "But, of course, I'm only claiming that you *morally* ought to give that much; I admit, though, that, all things considered, you have sufficient reason not to do so."