

CONSEQUENTIALISM AND COORDINATION PROBLEMS

Douglas W. Portmore

According to consequentialism, whether or not you should perform an act, ϕ , depends solely on the value associated with your ϕ -ing.¹ Yet, importantly, the value associated with your ϕ -ing can vary greatly depending on what other actions you and others will be simultaneously and subsequently performing. Imagine, for instance, that you're driving behind a slow-moving truck. Would it be good to accelerate? Well, it depends on whether you would simultaneously change lanes. For if you were to accelerate without changing lanes, you would crash into the back of the truck.²

Here, we have a coordination problem: the best available consequences will obtain only if you can coordinate your ϕ -ing (e.g., your using your foot to press down on the gas pedal to accelerate) with certain other actions (e.g., your using your hands to steer to the left to change lanes). More generally, the sort of coordination problem that I'm interested in has the following structure: (1) S_1 is deliberating at t about whether or not to ϕ at t' and (2) although S_1 's ϕ -ing at t' would not itself have good consequences, good consequences would ensue if both S_1 ϕ s at t' and S_2 ψ s at t'' , where S_1 may or may

* Check <https://dl.dropboxusercontent.com/u/14740340/Consequentialism%20and%20Coordination%20Problems.pdf> to be sure that you have the latest draft.

¹ Elsewhere I have used the term 'consequentialism' more broadly (see, for instance, Portmore 2011), but here I will use it to refer exclusively to those theories that hold that the morality of an agent's ϕ -ing depends solely on the impersonal (or agent-neutral) value associated with her ϕ -ing.

² This is a revised version of an example that comes from Goldman 1978, p. 186.

not be identical to S_2 and where $t < t' \leq t''$.³ In the above case, which I'll call *The Car*, S_1 and S_2 both refer to you, t' and t'' both refer to the time that's immediately subsequent to the present time t , ϕ -ing refers to accelerating, and ψ -ing refers to changing lanes.

In this paper, I will consider how consequentialists should treat S_2 and the possibility that she will ψ at t'' . At one end of the spectrum, consequentialists would hold that, in deciding whether or not to ϕ at t' , S_1 should always treat S_2 as a force of nature over which she has no control and, thus, treat the possibility that S_2 will ψ at t'' as she would the possibility that a hurricane will take a certain path. On this view, S_1 is to predict whether or not S_2 will ψ and act accordingly. At the other end of the spectrum, consequentialists would hold that S_1 should always treat S_2 as someone available for mutual cooperation and, thus, treat the possibility that S_2 will ψ at t'' as something to be relied upon. On this view, S_1 is to rely on S_2 's cooperation and so play her part in the best cooperative scheme involving the two of them. A third and intermediate position would be to hold that whether S_1 should treat S_2 as a force of nature or as someone available for mutual cooperation depends on whether S_1 can see to it that S_2 will ψ at t'' by, say, having the right set attitudes. I'll argue for this third position. As we'll see, an important implication of this view is that consequentialists should be concerned not just with an agent's voluntary actions but also with their involuntary acquisitions of various mental attitudes, such as beliefs, desires, and intentions. Indeed, I will argue that

³ Since S_1 may or may not be identical to S_2 , one possibility is that S_1 and S_2 are just two time-slices of the same person with S_1 being the present time-slice and S_2 being some future time-slice, but another possibility is that S_1 is the present time-slice of one person and S_2 is the future time-slice of some other (numerically distinct) person. Also note that, as I'll use them, the symbol ' $<$ ' stands for 'is earlier than' and the symbol ' \leq ' stands for 'is earlier than or identical to'.

consequentialists should hold both that (1) an agent's most fundamental duty is to have all those attitudes that she has decisive reason to have and only those attitudes that she has sufficient reason to have and that (2) she has a derivative duty to perform an act ϕ if and only if her fulfilling this fundamental duty ensures that she ϕ s.

1. Two Extreme Positions

As I noted, the position that I'll be arguing for lies between two extremes. At one end of the spectrum, we have *actualist consequentialism*, which holds that we should, in every instance, determine whether or not an agent should ϕ by looking at the actual consequences of her ϕ -ing. More precisely, the view holds:

AC A subject S is, as of time t , morally required to perform an act ϕ at t' if and only if S's ϕ -ing at t' is, as of t , an option for S and the prospect of S's ϕ -ing at t' is better than that of any alternative option that S has as of t .⁴

Now, if determinism is true, it's easy enough to talk about the consequences of an act, for there will be some determinate fact as to what they will be. But if indeterminism is true, then there will be no such determinate fact. Instead, there will be only a probability distribution over the possible worlds that may result from the

⁴ A subject's properties can change over time. For instance, a subject may in the past have had the property of being required to ϕ even though she doesn't now have this property due to her present inability to ϕ . And, on my parlance, {S is, as of time t , morally required to ϕ at t' } is equivalent to {S has at t the property of being morally required to ϕ at t' }. Likewise, {S's ϕ -ing at t' is, as of t , an option for S} is equivalent to {S has at t the property of having ϕ -ing at t' as an option}.

performance of that act. And since indeterminism may be true, I have chosen to formulate AC in terms prospects as opposed to determinate consequences. And given that we'll be talking about prospects, we'll need a way to evaluate them, and that is as follows: the prospect of ϕ is better than the prospect of ψ if and only if $\sum_i [Pr(w_i/\phi) \times V(w_i)]$ is greater than $\sum_i [Pr(w_i/\psi) \times V(w_i)]$, where w_i is a possible world, $Pr(w_i/\phi)$ is the objective probability of w_i given ϕ , and $V(w_i)$ is the value of w_i .⁵

To understand AC better, recall the sort of coordination problem that I'm interested in: (1) S_1 is deliberating at t about whether or not to ϕ at t' and (2) although S_1 's ϕ -ing at t' would not itself have good consequences, good consequences would ensue if both S_1 ϕ s at t' and S_2 ψ s at t'' , where S_1 may or may not be identical to S_2 and where $t < t' \leq t''$. On AC, S_1 is to treat the possibility that S_2 will ψ at t'' as she would the possibility that a hurricane will take a certain path. That is, she is to predict its course and act accordingly. This is true, on AC, regardless of whether S_1 is identical to S_2 , regardless of whether t' is identical t'' , and regardless of whether S_1 has, as of t , just as much control over whether or not S_2 will ψ at t'' as she has over whether or not she will ϕ at t' . To illustrate, consider the case from above and imagine that you are so frustrated with the slow-moving truck in front of you that you have decided to accelerate so as to crash into the back of the truck as hard as you can. Given this decision, you would not change lanes if you were to accelerate. So the prospect of your accelerating is much worse than the prospect of your not accelerating. Thus, AC implies that you should not

⁵ The objective probability of an event depends only on the facts and not on anyone's beliefs or evidence. If determinism is true, the objective probability of an event is always either zero or one.

accelerate. But AC has this implication even if your job depends on your accelerating while changing lanes. For suppose that the only way that you'll get to work on time—and let's assume that your job depends on your getting to work on time—is if you accelerate while changing lanes. If you accelerate without changing lanes, you will rear-end the truck, which will result in your losing both your job and your driver's license. If you change lanes without accelerating, you will be rear-ended by the Ferrari that's coming up fast from behind in the only adjacent lane, and this too will result in your losing both your job and your license. And, if you just stay behind the slow-moving truck at the same rate of speed, you'll be late for work, lose your job, but keep your license. So the only way to keep both your job and your license is to accelerate while changing lanes. Assume for the sake of simplicity that the relevant objective probabilities are all either zero or one. Thus, the probability that you'll be rear-ended by the Ferrari if you change lanes without accelerating is one, and the probability that you'll keep your job if you are late for work is zero. Given this, the relevant possibilities and their associated values can be summarized fairly simply, as they are in Table 1.

Table 1

<i>The Car</i>	Hand: Change Lanes	Hand: Don't change lanes
Foot: Accelerate	0	- 10
Foot: Don't accelerate	- 10	- 5

Actualist consequentialism (AC) holds that when deciding what you should do with your foot (such as whether to press down on the gas pedal so as to accelerate), you must treat what your hands will simultaneously do (say, steer straight ahead so as not to change lanes) as an immutable force of nature. You are to predict its course and act accordingly. And since your intention is to rear-end the slow-moving truck, you can confidently and accurately predict that your hands would not steer the wheel so as to change lanes even if you were to press down on the gas pedal with your foot. Thus, AC implies that you should not press down on the gas pedal with your foot even though you must do so in order to keep your job.⁶ For you must press down on the gas pedal in order to accelerate, and you must accelerate while changing lanes in order to keep your job.

The problem, then, with AC is that it has you deliberate as follows: “(1) I have a conditional obligation to refrain from accelerating if I’m not going to change lanes; (2) I’m not, as a matter of fact, going to change lanes (given my intention to rear-end the truck); and, therefore, (3) I have an unconditional obligation to refrain from accelerating.” The problem with deliberating in this way is that it treats the truth of (2) as fixed for the purposes of deliberation when its truth depends on the course of your

⁶ Although AC implies that it is impermissible for you to press down on the gas pedal with your foot to accelerate given that there is an alternative (such as, not pressing down on the gas pedal and just continuing to follow slowly behind the truck) whose prospect is better than its prospect, AC implies that you are required to perform the conjunctive act consisting of your pressing down on the gas pedal with your foot to accelerate while steering with your hand to change lanes, for the prospect of your performing this conjunctive act is better than that of any alternative option that you have. Thus, AC conflicts with the following principle of standard deontic logic: $R(S, t_i, [x_1, x_2, \dots, \& x_n]) \rightarrow [R(S, t_i, x_1), R(S, t_i, x_2), \dots, \& R(S, t_i, x_n)]$, where ‘ $R(S, t_i, \phi)$ ’ stands for ‘ S is, as of t_i , required to ϕ ’. Strangely, proponents of AC are content with this—see, for instance, Jackson and Pargetter 1986. And see Portmore 2011, pp. 181–183, for further discussion.

deliberations. For suppose that you were instead to deliberate as follows: “(i) I have an unconditional obligation to act so as to keep both my job and my driver’s license, (ii) in order to keep both my job and my driver’s license, I must accelerate while changing lanes, and, therefore, (iii) I have an unconditional obligation to accelerate while changing lanes.” If you were to deliberate in this way and, as a result, form the intention to accelerate while changing lanes, then this is precisely what you would do. And, if this is what you were to do, then (2) would be false.

So it seems that we should treat your hands’ steering straight ahead not as some immutable force of nature, but as an event over which you have control. For, indeed, you have as much control over what your hands will do as you have over what your foot will do. Both your hands and your feet will behave in accordance with your intentions, or so we’re assuming. So in deciding what to do with your foot (e.g., whether to press down on the pedal to accelerate), you should look not to what your hands would do given your present intention, but to what your hands could, and would, do if you formed the right intention. And, if you were to form the right intention—that is, the intention both to press down on the gas pedal with your foot to accelerate and to steer with your hands to change lanes, then that’s exactly what you would do.

This brings us to the other end of the spectrum: *possibilist consequentialism*, which holds that we should, in every instance, determine whether or not an agent should ϕ by looking at the possible consequences associated with her ϕ -ing, where the possible consequences are those possible given what other actions she and others could simultaneously and subsequently perform. More precisely, the view holds:

PC S is, as of t , morally required to ϕ at t' if and only if S's ϕ -ing at t' is part of every optimal maximal pattern of action that's possible as of t .

A *pattern* of action is a set of actions that spans across times and/or individuals. So my picking up takeout and then the two of us sitting down to eat in together is one pattern of action, and the two of us going out for dinner is another. And some patterns of action are more specific than others, where one pattern (MP₁) is more specific than another (MP₂) if and only if the performance of the one (MP₁) logically necessitates the performance of the other (MP₂) but not vice versa. Thus, my picking up takeout and then the two of us sitting down to eat in together is more specific than my picking up takeout. And my picking up takeout from the Chinese restaurant across the street is more specific than my picking up takeout. A *maximal* pattern action, then, is one that is maximally specific, where a possible pattern of action is a maximally specific if and only if there is no other possible pattern of action that is more specific than it. And a pattern of action is *possible* if and only if all the acts in which that pattern consists would be performed provided that all the agents involved in that pattern had the right intentions at the right times.⁷ More precisely, a pattern of action is possible as of t if and only if there is some schedule of intentions, I , extending over both a set of agents, S , and a time-interval, T , beginning at t' (where t' is just subsequent t) such that the following are all true: (a) if

⁷ Here, by 'the right intentions', I don't mean 'the correct or appropriate intentions', but instead 'the specific intentions that are required to realize the pattern'.

the intentions of those in S were to follow I , then all the acts in which that pattern consists would be performed; (b) the agents in S have at t the capacity to continue, or to come, to have the intentions that I specifies for t' ; and (c) for any time t'' in T after t' ($t' < t''$), if the intentions of those in S followed I up until t'' , then those in S would have just before t'' the capacity to continue, or to come, to have the intentions that I specifies for t'' .⁸ Lastly, a maximal pattern of action that's possible as of t is an *optimal* one if and only if there is no other maximal pattern of action that's possible as of t whose prospect is better than its prospect.

The upshot of this definition of {possible as of t } with respect to a pattern of action is that a pattern of action can be possible even if that pattern will not be realized, and even if no single agent can, at t , see to it that it will be realized. Thus, the pattern of action consisting of your accelerating while changing lanes is a possible one even though it will not be realized. It will not be realized, because you intend to accelerate while steering straight ahead and to thereby rear-end the truck, and because you will succeed in doing as you intend. Moreover, the pattern of action consisting of my picking-up takeout and then the two of us sitting down to eat in together is a possible pattern of action even if we hate each other so much that neither of us would ever agree to sit down to eat with the other no matter what anyone says or does. This pattern of action is, nevertheless, possible, because there is a schedule of intentions that meets criteria (a)–(c) above. That is, if, after I return with the food, each of us were to form the intention to sit

⁸ This definition of what it is for a pattern of action to be possible as of t is modeled after Ross's definition of a performable option—see Ross 2012, p. 81.

down to eat together, then that's what we would do. There may be nothing that either of us can do to see to it that the other forms this intention, but so long as we each have the capacity to form the requisite intentions at the right times and would execute every act in the pattern if we did, then the pattern is possible in the sense defined above.

Unlike actualist consequentialism, possibilist consequentialism has plausible implications in *The Car*. Possibilist consequentialism implies that you're morally required to accelerate.⁹ After all, it's clear that all the best maximal patterns of action that are possible as of *t* include your accelerating while changing lanes so that you can keep both your job and your driver's license. Unfortunately, though, possibilist consequentialism has implausible implications in the following case. I call it *The Buttons*, because it involves two individuals, Coop and Uncoop, each with a button in front of them. Depending on whether each pushes their button, the consequences will be as depicted in Table 2.

Table 2

<i>The Buttons</i>	Uncoop: Push	Uncoop: Not-push
Coop: Push	10	0
Coop: Not-push	0	6

⁹ If you're having trouble seeing how you could be *morally* required to accelerate, then just imagine that you have a moral obligation to take care of your family and that to do so you must hold onto both your job and your driver's license.

Assume that Coop desires that the two of them cooperate—that is, he desires that they work together, coordinating their behavior with each other, so that they bring about the best possible outcome, the one in which they both push and, thereby, produce ten units of value. But, unfortunately, Uncoop doesn't desire this. And Coop is no dupe; he will push his button if and only if he knows that Uncoop desires to cooperate—that is, desires the state of affairs in which they work together to produce the optimal ten units of value. And let's assume that it is transparent to each of them whether the other desires this state of affairs or not. For let's suppose that they both belong to a strange breed of hominids that are like us in every way except that a 'C' appears on their foreheads when, and only when, they desire that everyone cooperates with each.¹⁰ So Uncoop will not push, because he doesn't want there to be cooperation. And Coop will not push because he knows that Uncoop is unwilling to cooperate, as no 'C' appears on his forehead. Indeed, Uncoop is so steadfast in his unwillingness to cooperate that he won't push regardless what anyone does, thinks, or feels.

So, as I've noted, Coop will not push, but we can still ask: Should he? And, according to possibilist consequentialism (PC), the answer is that Coop should. Indeed, PC implies that Coop is required to push his button. For Coop's pushing his button is part of every optimal maximal pattern of action that is possible as of t . All the optimal maximal patterns of action that are possible as of t include both Coop's pushing his button at t' and Uncoop's pushing his button at t'' (here, $t' = t''$), for only if they both

¹⁰ I borrow the idea of indicating the presence of a desire to cooperate with letters on their foreheads from Richard Yetter Chappell—see Chappell 2012.

push will the optimum ten units of value be achieved. If neither pushes, a mere six units of value will be achieved. And, if only one pushes, then zero units of value will be achieved. So if Coop abides by PC, he'll push his button, and this will result in zero as opposed to six units of value being achieved. But requiring Coop to push in this situation seems absurd. If Coop knows that he is unable to influence Uncoop's behavior, then Coop should treat Uncoop's behavior as a force of nature, like a hurricane. That is, he should predict its course and act so as to secure the best possible consequences *given its predicted course*.

To illustrate the absurdity of PC, consider the following analogy. Imagine that there are several paths that a hurricane could take depending on what exactly the antecedent conditions and the laws of nature happen to be. And it is because we are uncertain as to what exactly these conditions and laws are that we are uncertain as to what path the hurricane will take. Assume, then, that our uncertainty as to the path of the hurricane is due entirely to our ignorance of the relevant conditions and laws and not at all to any indeterminacy in the natural world. Thus, there is a determinate fact as to what path the hurricane will take. Let's assume that it will take what I'll call 'path 2'. Lastly, assume that depending on which path the hurricane takes and on which evacuation route you take, the consequences will be as depicted in Table 3.

Table 3

<i>The Evacuation</i>	The Hurricane: Takes path 1	The Hurricane: Takes path 2
You: Evacuate via route 1	2	- 10

You: Evacuate via route 2	- 10	1
----------------------------------	------	---

Now, there is a sense in which the antecedent conditions and/or the laws of nature could have been different, such that the hurricane would have been determined to take path 1 instead. So, in some sense, it is possible for the hurricane to take path 1. It's not physically possible given the actual conditions and laws. But had the conditions and/or the laws been different (and, in some sense, they could have been), the hurricane would have taken that path. Indeed, the sense of possibility here seems no different from the one in which it is possible that Coop and Uncoop both push despite the fact that Uncoop will not, in fact, push. But although there is indeed some sense in which it is possible for the hurricane to take path 1, I can't imagine that anyone thinks that you should, given this, evacuate via route 1. If the hurricane *will* take path 2 and you are powerless to change this, then you ought to evacuate via route 2. Likewise, in *The Buttons*, we should say that given both that Uncoop *will not* push and that Coop is powerless to change this, Coop should refrain from pushing as well. The fact that Uncoop *could* push in the sense that he would push if things were different than they are just seems irrelevant to Coop's decision.

PC treats these two cases (*The Buttons* and *The Evacuation*) differently only because it makes an arbitrary distinction between acts of nature and acts of agents. But arguable what's relevant is not whether the actions are those of an agent or of, say, some hurricane, but whether one has any control over how the given entity (agent or hurricane) will behave.

The lesson in all this, I believe, is that we need to find some consequentialist view that lies between these two extremes. At one extreme, actualist consequentialism (AC) holds that, in every instance of the coordination problem that I've been describing, S_1 is to treat the possibility that S_2 will ψ at t'' as she would the possibility that the hurricane in *The Evacuation* will take path 1. But we should reject this view, for, in some instances (such as in *The Car*), S_1 has, at t , as much control over whether or not S_2 will ψ at t'' as she does over whether or not she will ϕ at t' . And S_1 should not treat S_2 as a force of nature if she has as much control over whether or not S_2 will ψ at t'' as she does over whether or not she will ϕ at t' . At the other extreme, possibilist consequentialism (PC) holds that, in every instance, S_1 is to treat the possibility that S_2 will ψ at t'' as she would the possibility that she will change lanes in *The Car*. But we should reject this view as well, for, in some instances, S_1 has no more control over whether or not S_2 will ψ at t'' than she does over what path the hurricane will take in *The Evacuation*. So what we need is a version of consequentialism that is sensitive to whether or not S_1 can see to it that S_2 will ψ at t'' . It is to these sorts of views that I now turn.

2. Cooperative Consequentialism

Donald Regan's *cooperative consequentialism* is one view that lies between these two extremes, but, as we'll see, it's untenable.¹¹ Even so, it is well worth considering, for it

¹¹ Regan calls his theory cooperative *utilitarianism* as opposed to cooperative *consequentialism*, but given both that he is not concerned to advance any particular view (let alone a utilitarian one) about what consequences are good and that he explicitly says that he chose the former term only for the sake of familiarity, it seems that the latter term is actually the more appropriate one. See Regan 1980, p. 1.

teaches us some important lessons, which will, in the end, help us to find a better theory.

In any case, the view holds:

CC S is, at all times, required to correctly *apply* the CC decision procedure.

And S is, as of t , required *perform* an act ϕ at t' if and only if this is what S would do if S were to correctly apply the CC decision procedure.

Cooperative consequentialism is unlike most moral theories (including AC and PC) in that it requires agents to do more than just *perform* certain voluntary actions. It requires that they also correctly *apply* a certain decision procedure: viz., the CC decision procedure. The CC decision procedure is quite complex, but the basic idea, which is sufficient for our purposes, is simple enough: each agent should (1) hold "himself ready to do his part in the best pattern of behavior for the group of cooperators," (2) "identify the other agents who are willing and able to co-operate in the production of the best possible consequences," and (3) "do his part in the best plan of behaviour for the group consisting of himself and the others so identified, in view of the behaviour of non-members of that group" (Regan 1980, pp. x and 135). So, on CC, "correct behaviour is required, but certain attitudes and beliefs are required as well" (1980, p. 124). Agents must have a desire or willingness to cooperate with each other and true beliefs about which others are willing to cooperate. Thus, as Regan puts it, CC is not *exclusively act-*

orientated.¹² If an agent performs the act that she would have been directed to perform had she correctly applied CC's decision procedure but does so without correctly applying CC's decision procedure, then although she has performed the right act, she has not done all that the theory requires of her.

To illustrate this unique and important feature of CC, consider a variant on *The Buttons* in which Coop does not hold himself ready to cooperate. Assume that he doesn't even bother to consider the lack of a 'C' on Uncoop's forehead, forms no belief as to whether Uncoop is willing to cooperate, and not-pushes simply as a result of inertia. If so, Coop has performed the right act, for given Uncoop's unwillingness to cooperate, the best plan of behavior for the group consisting of himself and all those willing to cooperate (the empty set) is the one where he not-pushes. But he has not done all that CC requires of him. For, on CC, Coop was required not only to not-push but also to identify Uncoop as a non-cooperator and to believe him to be such, which is something that he failed to do.

Requiring of agents something more than just the performance certain acts may seem unusual for a moral theory, but it is essential to CC's ability to avoid the counter-intuitive implications that plague actualist consequentialism (AC).¹³ Besides AC's

¹² A theory T is not exclusively act-orientated if and only if T requires of agents something more than just the performance of certain voluntary acts. This is not Regan's definition, for he provides no definition—see 1980, p. 109. But I believe that this definition captures (at least, roughly and sufficiently well for our purposes) the notion that he has in mind.

¹³ Perhaps, it does not seem terribly unusual. There are, after all, moral theories (such as global consequentialism—see Pettit and Smith 2000) that require agents to have certain motives, intentions, and character traits, and not just to perform whatever acts might result in their acquiring them. Nevertheless, what makes CC stand out as being particularly unusual is that it makes an agent's obligation to perform an act derivative of her obligation to do that something more. Thus, on CC, an agent's fundamental obligation

counter-intuitive implications in *The Car*, AC runs into trouble with *The Buttons*—that is, with the original version of the case, the one in which Coop desires the state of affairs in which they cooperate with each other. The problem with AC is that it fails to require Uncoop to see to it that the best possible outcome is realized. And all he would need to do to see to it that the best possible outcome is realized is to form both a desire for mutual cooperation and an intention to push his button. The latter will ensure that he pushes, and the former will ensure that Coop pushes. For if Uncoop forms a desire for mutual cooperation, then a ‘C’ will appear on his forehead. And if a ‘C’ appears on Uncoop’s forehead, Coop will push. Thus, if Uncoop forms both a desire for mutual cooperation and an intention to push his own button, then both he and Coop will both push, producing the optimal ten units of value. Yet, AC doesn’t require Uncoop to form a desire for mutual cooperation. Nor does it require Uncoop to push. Strangely, AC requires Uncoop to not-push, because given Uncoop’s lack of a desire for mutual cooperation and, consequent, lack of a ‘C’ appearing on his forehead, Coop is not going to push. And given that Coop is not going to push, the prospect of Uncoop’s not-pushing is better than the prospect of Uncoop’s pushing. Thus, Uncoop should, according to AC, not-push.

These problems arise for AC because it is focused exclusively on acts when whether the best possible outcome will be realized depends not only on whether Uncoop will push, but also on whether Uncoop desire that he and Coop cooperate. But

is to correctly apply the CC decision procedure, and she has a derivative obligation to perform a certain act if and only if she would perform that act if she were to fulfill this fundamental obligation.

being exclusively act-orientated, AC is unable to require that Uncoop have any desire. Of course, AC could require Uncoop to perform whatever voluntary act might cause him to desire their mutual cooperation—perhaps, the act of going to a talented hypnotist who can implant the desire in him. But let's assume that there is no such available act and, thus, that the only way that Uncoop can come to desire their mutual cooperation is by recognizing and responding appropriately to the decisive reason that he has to desire their mutual cooperation—viz., that their mutual cooperation will produce the most good.¹⁴ This, I take it, is the typical way in which agents come to have attitudes such as beliefs, desires, and intentions. Typically, they form these attitudes not by intending to form them, but by responding to reasons for having them in an involuntary and non-agential way—the way that one comes to believe that the floor is wet when walks on a wet floor with one's bare feet.

The lesson, here, is that if we're going to come up with a theory that avoids AC's troubles with *The Buttons*, it will need to be one that's not exclusively act-orientated. CC avoids being exclusively act-orientated by requiring agents not only to perform certain actions but also to correctly apply the CC decision procedure. Unfortunately, though, its CC's requiring agents to correctly apply the CC decision procedure that is also its downfall. For there are two devastating objections to CC's doing so. The first is that, in doing so, CC runs afoul of the 'ought'-implies-'can' principle, because to require agents

¹⁴ I'm assuming that the fact that their cooperating would produce the most good constitutes a decisive reason for Uncoop to desire their mutual cooperation. Thus, it is a fact about the object of the desire, not any fact about the having of the desire, that constitutes a decisive reason for Uncoop to desire their mutual cooperation. Thus, I avoid any worries associated with the wrong-kind-of-reasons problem. See Rabinowicz and Rønnow-Rasmussen 2004 for a discussion of the problem.

to correctly apply the CC decision procedure is to require them to form certain beliefs and attitudes whether they have the capacity to form them or not.¹⁵ For instance, CC requires agents to form true beliefs about who is and isn't willing to cooperate, about what all the non-cooperators are going to do, and about what the best pattern of behavior for the group of cooperators is given what all the non-cooperators are going to do (Regan 1980, pp. 151–158). But an agent may not have the mental capacity to form such beliefs. Suppose, for instance, that the relevant agent is an infant boy who has the capacity to form the intention to push a button, but not the capacity to form a belief about what the best pattern of behavior is. Perhaps, he lacks the relevant concepts to even conceptualize such a belief. Or, perhaps, a mad neuroscientist has inserted electrodes into his brain so as to make it impossible for him to form such a belief (Barley 1984, p. 158). Regardless, CC requires the boy to form a true belief about what the best pattern of behavior is. And this is absurd. The boy can't be required to form beliefs that he lacks the capacity to form.

The second problem with CC's requiring agents to apply the CC decision procedure is that sometimes the costs associated with applying the CC decision procedure far outweigh any of its benefits. And, in some instances, applying the CC decision procedure may even be disastrous. And no consequentialist theory should require an agent to ϕ (e.g., to apply the CC-decision procedure) if whether or not she ϕ s is a voluntary choice and choosing to ϕ would have disastrous consequences.¹⁶ To

¹⁵ This objection originates with Barley 1984, pp. 158–159.

¹⁶ Regan does talk of there being a *choice* between applying the CC decision procedure and applying some other decision procedure—see 1980, p. 179.

illustrate, suppose that there is a mad telepath who will monitor Coop's decision-making process in *The Buttons* and will blow up Macy's, or perhaps even the whole world, if Coop applies the CC decision procedure. And assume that whether or not Coop applies the CC decision procedure is a voluntary choice. I'm assuming that although fully and correctly applying the CC decision procedure will involve more than just performing certain voluntary acts (it will, for instance, involve having certain beliefs and other attitudes), choosing to proceed via this decision procedure is a voluntary choice. In any case, let's call this revised version of *The Buttons: The Mad Telepath*. Even in cases such as *The Mad Telepath*, CC requires agents to apply the CC decision procedure (Regan 1980, p. 177). This is problematic for CC, for besides its being counterintuitive to think that Coop is required to voluntarily choose to apply the CC decision procedure when doing so would be so disastrous, this implication of CC seems anti-consequentialist in spirit. Thus, we might wonder not only whether CC is a *plausible* consequentialist theory, but also whether it is even a *consequentialist* theory at all.

Regan is well aware that CC has this sort of counterintuitive implication. Indeed, I borrow the mad telepath example (slightly revised) from him—see 1980, p. 177. But Regan is willing to stomach such implications, because he thinks both (1) that no theory that is exclusively act-orientated can avoid counterintuitive implications in cases like *The Buttons* and (2) that no theory that is not exclusively act-orientated can avoid counterintuitive implications in cases like *The Mad Telepath*.¹⁷ Now Regan would seem to

¹⁷ Another reason that he is willing to accept CC despite its counterintuitive implications is that CC is *adaptable*, where "a theory T is adaptable if and only if the agents who satisfy T, whoever and however numerous they may be, are guaranteed to produce the best [act-]consequences possible as a group, given the

be correct about (1). No exclusively act-orientated theory can avoid counterintuitive implications in cases like *The Buttons*. In such cases, the relevant agent can see to it that the best possible outcome is realized only by having certain attitudes. Performing certain actions is not enough. To illustrate, consider yet another variant on *The Buttons*. Assume that, in this variant, Uncoop pushes his button but without desiring their mutual cooperation. Assume that he pushes only because his commanding officer orders him to do so. Given that Uncoop lacks a desire for their mutual cooperation, no 'C' appears on his forehead and, consequently, Coop does not push. Thus, the best possible outcome is left unrealized despite the fact that Uncoop pushes. Clearly, then, it is not always sufficient to perform certain acts. Sometimes, the only way to ensure that the best possible outcome is realized is to have certain attitudes.

But even if Regan is correct about (1), he seems to be incorrect about (2). Regan defends (2) by arguing that since any theory that's not exclusively act-orientated will have to require 'something more' of agents than the mere performance of certain

behavior of everyone else" (Regan 1980, p. 6). The reason that I've added the modifier 'act-' to 'consequences' is that Regan specifically tells us that he's not interested in all the consequences associated with satisfying T, but only with the consequences that stem from performing the acts that satisfying T entails performing. So if, to satisfy T, an agent must not only perform certain acts but also apply a certain decision procedure (as is the case with CC), then in determining whether CC is adaptable or not we are to exclude from consideration "any consequences of the application of [the]...required decision procedure *aside from* the consequences of the act or acts chosen by that procedure" (1980, p. 8). Thus, in determining the best act-consequences possible in *The Mad Telepath* case, we are to ignore the destruction of Macy's and focus solely on the consequences of the agent's pushing or not-pushing. Now, I see why it would be desirable for a theory to be *truly adaptable*, where a theory is truly adaptable if and only if the agents who satisfy T, whoever and however numerous they may be are guaranteed to produce the best consequences (and not merely the best act-consequences) possible as a group, given the behavior of everyone else. But I fail to see why it is desirable for a theory to be adaptable in Regan's sense. And, unfortunately, a theory that is truly adaptable—"a theory which guarantees in all cases that the best possible consequences *all things considered* will be produced by whatever collection of agents satisfy the theory—is [as Regan shows] a logical impossibility" (1980, p. 184).

voluntary acts, “there is always the possibility that there will be a mad telepath in some situation who will blow up Macy’s [or the whole world] in response to that ‘something more’” (Regan 1980, p. 181). But, from this, it doesn’t follow that every theory that is not exclusively act-orientated must require agents to voluntarily choose to do that ‘something more’ even when the bad consequences of that something more by far outweigh whatever good consequences it may have, which is exactly what gets CC in trouble. For instance, a theory that required Coop to decide to not-push by way of choosing to apply a certain decision procedure (say, CC’s decision procedure) but only when the good effects of choosing to apply this procedure outweighed the bad effects of doing so would not be exclusively act-orientated and, yet, unlike CC, would avoid counterintuitive implications in *The Mad Telepath*.¹⁸

So, as I said at the beginning of this section, CC is untenable. It violates the ‘ought’-implies-‘can’ principle and has some very counterintuitive implications. But it was, as I promised, well worth considering, for, in doing so, we have learned some important lessons—three lessons, in fact: (L1) a moral theory should not be exclusively act-orientated—more specifically, it should require agents to have certain attitudes and not just to perform certain voluntary actions; (L2) even if a moral theory is not exclusively act-orientated, it should never require an agent to voluntarily choose to ϕ when ϕ -ing would have disastrous (or even suboptimal) consequences; and (L3) a moral theory should avoid running afoul of the ‘ought’-implies-‘can’ principle and so it should

¹⁸ I’m not claiming that this revised version of CC doesn’t have its own problems.

require agents to form and perform only those attitudes and actions that they can form and perform.

In addition to these three lessons, we have, from section 1, learned a fourth: (L4) a moral theory should lie somewhere between the extremes of actualist consequentialism and possibilist consequentialism and, thus, should be sensitive to whether or not the agent can, in sorts of coordination cases that I've been discussing, see to it that relevant coordination takes place—that is, a moral theory should treat *other acts*¹⁹ (whether they be concurrent or subsequent and whether they be performed by the agent herself or by others) as forces of nature only when the agent has no more control over whether these other acts will be performed than she does over whether a hurricane will take a certain path.

In the next section, I make use of these lessons in formulating a more plausible alternative to cooperative consequentialism.

3. Attitude-consequentialism

One way to take all these lessons to heart is to adopt the view that I call *attitude-consequentialism*. On attitude-consequentialism, agents are, first and foremost, required to have a permissible set of attitudes—that is, they're required to have some set that includes all those attitudes that they have decisive reason to have and only those attitudes that they have sufficient reason to have. Now, although attitude-

¹⁹ If we're assessing the deontic status of *S*'s ϕ -ing at *t* and/or *S* is deliberating about whether or not to ϕ at *t*, then the other acts are all and only those that are not *S*'s ϕ -ing at *t*.

consequentialism is neutral both with regard to what agents are permitted to desire and with regard to what agents are permitted to believe, it is not neutral with regard to all attitudes. Given that it is a moral theory and, thus, must ultimately provide an account of which voluntary acts are permissible, and given that it is only through intending to perform some act that an agent can voluntarily perform it, it must offer us some account of what agents are permitted to intend to do. And since it is a type of consequentialism, the account must be one that appeals to the consequences or prospects of the acts that these agents might intend to perform. Roughly, then, attitude-consequentialism holds that an agent is permitted to intend to perform an act ϕ if and only if there is no alternative act ψ such that the prospect of her ψ -ing while having permissible background attitudes is better than the prospect of her ϕ -ing while having permissible background attitudes. The *background* attitudes with respect to a given act are all the attitudes that the agent has besides the intention to perform that act.

You might wonder why I talk about background attitudes at all. Why not just hold that an agent is permitted to intend to perform an act ϕ if and only if there is no alternative act ψ such that the prospect of her ψ -ing is better than the prospect of her ϕ -ing? There are two reasons. First, we're ultimately concerned with whether a *set* of attitudes is permissible, and it can be permissible to have an attitude in combination with some attitudes but not in combination with other attitudes. For instance, I can permissibly intend not to leave the house today, but not if I intend to meet with a student at my office this morning and believe that my leaving the house is a necessary

means to my doing so. So we can't determine which intentions are permissible without considering the subject's background attitudes.

Second, we need to look at background attitudes, because the consequences of performing an act can vary greatly depending on what background attitudes the agent has. Consider again *The Buttons*. If Uncoop pushes without desiring that he and Coop cooperate, the consequences will be bad. But if Uncoop both pushes and desires that he and Coop cooperate, the consequences will be good. Likewise, if you, in *The Car*, accelerate without intending to change lanes, the consequences will be bad. But, if you accelerate while also intending to change lanes, the consequences will be good. So if we're looking for a consequentialist account of which acts agents should intend to perform, then, given that the consequences of those acts vary greatly depending on the agent's background attitudes, we must take those background attitudes into account.

I have yet to say how we are to determine which actions are required on attitude-consequentialism. The key to figuring out what attitude-consequentialism should, and in fact does, say about actions is to note that what an agent will do often depends on what attitudes she has. Oftentimes, if she intends to ϕ , she will ϕ ; but if she doesn't, she won't. It makes sense, then, for attitude-consequentialism to hold that an agent is required to ϕ if and only if she would ϕ so long as she were to have some permissible set of attitudes.

To understand attitude-consequentialism better, let's consider some of its implications in *The Buttons*. Suppose, for instance, that we want to know whether Uncoop is required to push. On attitude-consequentialism, whether he is or not depends

on what attitudes he's required to have and whether his having these (the required) attitudes guarantees that he ϕ s. So what are the attitudes that he is required to have? Well, one thing that I've been assuming is that he's required to desire that he and Coop cooperate. He's required to have this desire, not because having it would have good consequences, but because the object of this attitude (i.e., their mutual cooperation) has good consequences. The fact that an event E would have good consequences is an excellent reason to desire that E occurs. And since he has no better reason to desire that any alternative event occurs (such as the event of their mutual *non*-cooperation), we can conclude that Uncoop is required to desire that he and Coop cooperate. And given the description of the case, this is the only relevant background attitude that he's required to have.

Next, we are to figure out whether the prospect of Uncoop's both pushing and having this required desire is better than the prospect of Uncoop's both not-pushing and having this desire. And it is, for the prospect of the former is 10 units of value (the result of their both pushing), whereas the prospect of the latter is zero units of value (the result of Uncoop's not-pushing while Coop pushes). Therefore, attitude-consequentialism implies that Uncoop is required to intend to push his button. And if, as I'm assuming, Uncoop would push if he were both to intend to push and to desire that he and Coop cooperate, attitude-consequentialism further implies that Uncoop is required to push. So attitude-consequentialism requires Uncoop both to push and to desire he and Coop cooperate, thereby securing the best possible state of affairs: the one in which they both push and, thus, produce ten units of value.

To make sure that we have at least a basic understanding of attitude-consequentialism, let's look at its implications in *The Car* as well. The question, here, is whether you ought to accelerate. And to answer this question we must consider whether there is any alternative to your accelerating (such as your slowing down or your maintaining your current speed) such that the prospect of your intending to perform that alternative while having permissible background attitudes is better than the prospect your accelerating while having permissible background attitudes. And, in *The Car*, the only relevant background attitude with regard to your decision about whether or not to accelerate is your intention with regard to steering the vehicle. Now, since the only way for you to keep both your job and your driver's license is to accelerate while changing lanes, it seems that you ought to intend to accelerate while changing lanes. And if you ought intend to accelerate while changing lanes, then you ought both to intend to accelerate and to intend to change lanes. Thus, you ought (i.e., are required) to intend to change lanes. And the best prospect associated with this required background attitude is the one in which you accelerate and have this required attitude. So, in *The Car*, attitude-consequentialism holds that you're to accelerate and to intend to change lanes as you do so, thereby securing the best available consequences, the ones where you keep both your job and your driver's license.

The preceding serves only as a quick gloss on the view. At this point, though, we should be ready for a more precise and formal statement of the view. More precisely, then, attitude-consequentialism (or @C) is the conjunction of the following three theses, which I label FUN, INT, and PER:

- @C (FUN) *S*'s most *fundamental* obligation, as of *t*, is to have at *t'* some permissible set of attitudes.
- (INT) *S* is permitted, as of *t*, to *intend* at *t'* to ϕ at *t''* if and only if both
- (1) *S* has, as of *t*, the capacity to intend at *t'* to ϕ at *t''* and
 - (2) there is no alternative act ψ such that both
 - (a) *S* has, as of *t*, the capacity to intend at *t'* to ψ at *t''* and
 - (b) the prospect of *S*'s both ψ -ing at *t''* and having at *t'* permissible background attitudes is better than the prospect of *S*'s both ϕ -ing at *t''* and having at *t'* permissible background attitudes ($t < t' \leq t''$).
- (PER) *S* is, as of *t*, required to *perform* ϕ at *t''* if and only if *S* will ϕ at *t''* so long as she were to have at *t'* some permissible set of attitudes.

Although the above formulation is quite complex, the main idea is relatively simple: an agent's most fundamental obligation is to have some permissible set of attitudes, and she has a derivative obligation to perform an act if and only if she would perform that act so long as she were to have some permissible set of attitudes. The only additional element to the view is an account of the permissibility of *S*'s intending to ϕ according to which *S* is permitted to intend to ϕ so long as there is no alternative act ψ such that the prospect of *S*'s both ψ -ing and having permissible background attitudes is better than the prospect of *S*'s both ϕ -ing and having permissible background attitudes.

But, of course, there are some nuances to the above formulation that this simple

gloss papers over. So let me explain these nuances. First, you'll note that there are three time indices in the phrase {S is permitted, as of t , to *intend* at t' to ϕ at t'' }: there's the time at which S has the property of being permitted to have the intention (that is, t), there's the time at which the intention occurs (that is, t'), and there's the time at which the action that is intended occurs (that is, t''). The indices are necessary, for these could turn out to be three distinct times. For instance, I could *now* (at t) have the property of being permitted to intend *just subsequent from now* (at t') to mow the lawn *tomorrow* (at t'').

Second, the formulation @C includes two conditions not mentioned in my initial gloss: (1) and (a). The point of both is to avoid holding that an agent is required to form an attitude that she doesn't have the capacity to form, because we don't want to run afoul of the 'ought'-implies-'can' principle. Of course, it's difficult to say what exactly it means to have, as of t , the capacity to A at t' , where A is some attitude. But it is clear that we do sometimes hold people responsible for having or lacking certain attitudes—e.g., for having an unwarranted belief, for having certain racist attitudes, or for lacking a sufficient concern for the welfare of others. And whether we hold people responsible for having or lacking such attitudes depends on our thinking that they had the capacity to have different or additional attitudes. For instance, we wouldn't hold an unconscious patient responsible for not having an occurrent desire to live. Nor would we hold a small child responsible for not intending to develop the virtue of temperance when she is still too young to grasp the concepts needed to form such an intention. So I think that we can safely assume that in order for an agent to have a certain attitude, she must, at a

minimum, be conscious and have the necessary conceptual apparatus. There are, of course, other requirements, but that needn't concern us here. For our purposes, it's important only that we never require an agent to have an intention that they lack capacity for, and we do that by including conditions (1) and (a).

Third, the above formulation talks about *permissible* background attitudes. I've already explained the importance of talking about background attitudes, but I haven't explained why I restrict the relevant background attitudes to only those that are permissible. To understand the need for this restriction, note first that the attitudes in question (judgment-sensitive attitudes such as desires, beliefs, and intentions) are all of the kind that one acquires involuntarily in response to reasons.²⁰ For instance, I come to the belief that I should go to the emergency room not by intending to form this belief, but by reading online that I have all the symptoms of appendicitis and recognizing that this fact constitutes decisive reason for me to go to the emergency room. But if the attitudes in question are ones that an agent can acquire only involuntarily in the response to reasons, then it wouldn't make sense to require agents to come to have attitudes that they lack sufficient reason to have. For, in that case, they wouldn't be able to acquire them by responding appropriately to their reasons. Nor could they acquire them voluntarily by intending to form them given that they are not of the type that can be acquired by intending to form them. Thus, we should restrict the relevant background attitudes to those that are permissible—that is, to those for which there is

²⁰ Judgment-sensitive attitudes are those attitudes that are sensitive to one's judgments about reasons. Beliefs, desires, and intentions are judgment-sensitive, but hunger is not. See Scanlon 1998, pp. 18–20.

sufficient reason. Thus, in *The Buttons*, it's fair to require Uncoop to secure the world in which both parties push by requiring him both to intend to push and to desire their mutual cooperation, for Uncoop has sufficient reason both to intend to push and to desire to cooperate. But it wouldn't make sense to require Uncoop to believe something that he has no reason to believe even if his having this belief would allow Uncoop to secure an even better outcome. For given that he lacks sufficient reason to have this belief and so can't acquire it (at least, not directly) by responding to his reasons, it would be unreasonable for anyone to expect or require him to form it.

At this point, the reader should have a sufficient understanding of @C so that I can now proceed to demonstrate that @C takes to heart all of lessons L1–L4. We've already seen that @C takes account of L3. It avoids running afoul of the 'ought'-implies-'can' principle by requiring agents to have only those attitudes that they have the capacity for. This is the point of conditions (1) and (a). Moreover, it's clear that @C takes to heart L1 in that it is not an exclusively act-orientated theory. After all, it requires agents not only to perform certain actions (see PER) but also to have certain attitudes (see FUN). Thus, @C implies that, in *The Buttons*, that Uncoop is required not only to push but also to desire that he and Coop cooperate.

@C also takes L2 to heart, thereby avoiding counterintuitive implications in cases like *The Mad Telepath*. Although @C does require 'something more' of agents than just the performance of certain voluntary acts, it never requires an agent to voluntarily choose to ϕ when her ϕ -ing would have disastrous consequences. For instance, @C would never require an agent to voluntarily choose to apply a certain decision

procedure when doing so would have disastrous consequences. And, thus, @C avoids the sort of counterintuitive implications that plagued CC in *The Mad Telepath*. I must concede, though, that @C would require an agent to involuntarily form a certain attitude in response to her reasons if her reasons decisively favored that attitude. What's more, @C would require this even if there were some mad telepath who was standing by to destroy the world if she forms that attitude. But I don't think that this is problematic. We can't blame someone for involuntarily forming an attitude in response to the decisive reasons that she has to do so. And if there were some act that she could perform so as to block her involuntary formation of this attitude, @C would require her to perform it. This, then, would be a case where an agent is required to act so as to make herself the type of person who would form the attitudes that she ought not rationally to form. But I don't see anything problematic about this kind of "rational irrationality."²¹

Lastly, @C takes L4 to heart in that it lies between the extremes of actualist consequentialism (AC) and possibilist consequentialism (PC). Unlike AC, @C does require you to accelerate in *The Car*. Although this will have bad consequences if you maintain your present intention to smash into the back of the truck, @C requires you to change your intention, forming instead the intention to accelerate while changing lanes, thereby ensuring that you'll keep both your job and your driver's license. AC, unfortunately, is exclusively act-orientated. So it can't require you to change your intention. It can require only that you accelerate or that you not accelerate depending on what the consequences would be. And if your intention won't change (even though you

²¹ See Parfit 1984, §5.

could change it), then your accelerating would have disastrous consequences. But this is to overlook the fact that you have just as much control over whether you'll change lanes as you do over whether you'll accelerate. Thus, it seems foolish for you to hold fixed your not changing lanes when deliberating about whether or not to accelerate, as both equally depend on the course of your present deliberations.

And, unlike PC, @C doesn't implausibly require Coop to push even though he is powerless to get Uncoop to push and his pushing without Uncoop's pushing has the worst possible consequences. PC, like AC, makes the mistake of being insensitive to one's degree of control over whether the relevant coordination will take place. PC just takes it to the opposite extreme. Whereas AC treats as fixed actions that you have complete control over, PC treats as unfixed actions that you have absolutely no control over. Thus, PC implies that Coop should push even though Uncoop will not push and Coop is powerless to change this. But unlike PC, @C rightly requires Coop to refrain from pushing given both that Uncoop isn't going to push and that Coop is powerless to change this. Thus, @C takes L4 to heart by being sensitive to whether or not the agent can see to it that the relevant coordination will take place. If the agent can see to this, then the agent is required to do so whether it be through her own intentional actions or through her involuntarily forming certain attitudes in response to her reasons to do so. And if she cannot see to it that coordination will take place, then she is to secure the best possible world in which coordination doesn't take place.

4. Conclusion

As we've just seen @C takes into account all of lessons: L1–L4. I haven't argued that @C is the only theory that does so or that it is the best of those that do. But I don't yet see any better way to taking account of these four lessons. And it seems that @C gets intuitive results in all the cases that we've been considering. And I'm unaware of a consequentialist theory that does a better job in these or other cases. So it seems to me that @C is, at least, a promising consequentialist theory. And the hope is that even if it doesn't live up to its promise, I will have at least specified some important lessons that must be taken into account in seeking a replacement for it.²²

References

- Barley, B. (1984). "Review of *Utilitarianism and Cooperation* by Donald Regan," *Noûs* 18: 152–159.
- Chappell, R. Y. (2012). "Let's start with objective...." (5:32 PM, November 7, 2012). Comment to a blogpost: Richard Yetter Chappell's "Why if everyone did that?" Philosophy, et cetera. Retrieved August 12, 2013, from <URL><http://www.philosophyetc.net/2012/11/what-if-everyone-did-that.html?showComment=1352327576983#c5832896466393289467></URL>.
- Goldman, H. S. [now H. M. Smith]. (1978). "Doing the Best One Can." In A. I. Goldman and J. Kim (eds.), *Values and Morals*, pp. 185–214. Dordrecht: D. Reidel Publishing Company.

²² For helpful discussions concerning these issues, I thank Chrisoula Andreou and commentators on some of my posts at *PEA Soup* (<http://peasoup.typepad.com/peasoup/>).

- Jackson, F. and R. Pargetter (1986). "Oughts, Actions, and Actualism." *Philosophical Review* 95: 233–255.
- Parfit, D. (1984). *Reasons and Persons*. Oxford: Oxford University Press.
- Pettit, P. and M. Smith (2000). "Global Consequentialism." In B. Hooker, E. Mason, and D. E. Miller (eds.), *Morality, Rules and Consequences: A Critical Reader*, pp. 121–33, Edinburgh: Edinburgh University Press.
- Portmore, D. W. (2011). *Commonsense Consequentialism: Wherein morality meets rationality*. New York: Oxford University Press.
- Rabinowicz, W., and T. Rønnow-Rasmussen (2004). "The Strike of the Demon: On Fitting Pro-Attitudes and Value." *Ethics* 114: 391–423.
- Regan, D. (1980). *Utilitarianism and Co-operation*. New York: Oxford University Press.
- Ross, J. (2012). "Actualism, Possibilism, and Beyond." In M. Timmons (ed.) *Oxford Studies in Normative Ethics: Volume 2*, pp. 74–96, Oxford: Oxford University Press.
- Scanlon, T. M. (1998). *What We Owe to Each Other*. Cambridge, Mass.: Belknap Press.