

ACTS, ATTITUDES, AND RATIONAL CHOICE

Douglas W. Portmore

Abstract: In this paper, I argue that we have obligations not only to perform certain actions, but also to have certain attitudes (such as desires, beliefs, and intentions), and this despite the fact that we rarely, if ever, have direct voluntary control over our attitudes. Moreover, I argue that whatever obligations we have with respect to our actions derive from our obligations with respect to our attitudes. More specifically, I argue that an agent is obligated to perform an action if and only if it's the action that she would perform if she were to have the attitudes that she is required to have. This view, which I call *attitudism*, has three important implications. First, it implies that an adequate practical theory must not be exclusively act-orientated. That is, it must require more from us than just the performance of certain voluntary acts. Second, it implies that an adequate practical theory must be attitude-dependent. That is, it must hold that what we ought to do depends on what attitudes we ought to have. Third, it implies that no adequate practical theory can require us to perform acts that we would not perform even if we were to have the attitudes that we are required to have. I then show how these implications can help us both to address certain puzzling cases of rational choice and to understand why most typical practical theories (utilitarianism, virtue ethics, rational egoism, Rossian deontology, etc.) are mistaken.

It seems that we have obligations not only to perform certain actions, but also to have and to form certain attitudes. After all, we make claims such as: (i) you should feel ashamed of yourself; (ii) Christians are obliged to love their enemies; (iii) there is nothing wrong with wanting to win; (iv) you should be proud of what you've accomplished; (v) it's wrong to hate people because of their race; and (vi) given all the scientific evidence to the contrary, you shouldn't believe that humans and dinosaurs ever coexisted. Of course, some will object that we can't be obligated to ϕ unless our ϕ -ing is under our voluntary control and that, unlike actions, attitudes are not under our voluntary control. We cannot, for instance, form a belief at will. No matter how much I want to believe that I have an immortal soul, I cannot, given the lack of evidence, bring myself to believe this—at least, not by willing myself to believe it.

I'll address this objection shortly. But, for now, note that, by 'attitudes', I mean to refer to all and only those mental states that are responsive to a rational subject's awareness of reason-constituting facts.¹ Such mental states include intentions, propositional attitudes

* The latest draft can always be found at <http://tinyurl.com/mrckwmb>.

such as beliefs and desires, as well as certain emotions, such as hope, fear, admiration, contempt, guilt, and indignation. These are mental states that a rational subject will tend to have, or tend not to have, in response to reasons (or apparent reasons)—facts (or what are taken to be facts) that count for or against the attitudes in question. For instance, a rational subject will tend to believe that it will rain today in response to her awareness of a fact that constitutes conclusive reason for her having this belief, such as the fact that a reliable weather service has predicted that it will rain today. This conception of ‘attitudes’ includes many mental states, but excludes feelings of hunger, nausea, and dizziness, which are not responsive to a subject’s awareness of reason-constituting facts. Suppose, for instance, that I have too quickly consumed a good-sized meal and am still feeling hungry, as there has not yet been sufficient time for brain to receive the relevant physiological signals from my stomach. Even if I am aware that I’ve eaten more than enough to be satiated, my hunger is not responsive to this awareness. Instead, it is responsive only to the physiological signals that supposedly take about twenty minutes to travel from the stomach to the brain. So it is only those mental states that are responsive to reasons that count as attitudes. And this is important, because insofar as we think that our obligations are a function of our reasons, the fact that our attitudes, as much as our actions, are responsive to reasons would suggest that we can have obligations with respect to them as well.

Presuming that I’m right about our having obligations with respect to both acts and attitudes, this raises the question of what, if any, relationship there is between the two. That is, what is the relationship between the acts that we ought to perform and the attitudes that we ought to have? One possibility, *the acts-first view*, is that the attitudes that we ought to have are a function of the acts that we ought to perform. Perhaps, we ought to have the attitudes that would lead us to perform the acts that we ought to perform. Another possibility, *the neither-first view*, is that the two are completely unrelated or that each is a function of some third thing, such as how we ought to be. Suppose, for instance,

¹ What I’m calling *attitudes* is similar to what Scanlon calls *judgment-sensitive attitudes*: that is, attitudes that are sensitive to our judgments about reasons (1998, p. 20). But, like Parfit, I think that we can respond to reasons without having any judgments about what our reasons are. For, as Parfit notes, “we respond to reasons when we are aware of facts that give us these reasons, and this awareness leads us to believe, or want, or do what these facts give us reasons to believe, or want, or do” (2011, p. 493). Thus, we can respond to reasons without knowing that this is what we are doing (2011, p. 461).

that we ought to be virtuous. On one version of the neither-first view, then, we ought to have the attitudes and perform the acts that a virtuous person would have and perform. The third and final possibility, *the attitudes-first view*, is that the acts that we ought to perform are a function of the attitudes that we ought to have. I'll argue for this third possibility. Specifically, I'll argue for a particular version of the attitudes-first view, which I'll call *attitudism*.

The basic idea underlying attitudism is quite simple. An agent's obligation to perform an action derives from a more fundamental obligation to refrain from having an impermissible set of attitudes. Specifically, attitudism holds that an agent is obligated to perform an act ϕ if and only if, and because, she would ϕ provided she were to have some permissible set of attitudes. To illustrate, suppose that Kwame, a middle-aged man with heart disease, is currently experiencing severe chest pain and is obligated to have the following set of attitudes: (1) the belief that taking a nitroglycerin pill is necessary to prevent himself from dying of a heart attack (because, we'll assume, he possesses conclusive evidence for this belief), (2) the desire not to die of a heart attack (because, we'll assume, it would be desirable for him to survive), and (3) the intention to take a nitroglycerin pill (because performing this action would produce what's desirable). Further suppose that if he were to have these attitudes, he would take a nitroglycerin pill, for he has several such pills on hand for exactly this sort of situation. Attitudism implies, then, that he has a fundamental obligation to have attitudes 1–3 and a derivative obligation to take a nitroglycerin pill, for this is the act that he would perform if he were to have the attitudes that he is required to have.

So that's the basic idea. More formally, though, the view is as follows.

Attitudism: For any rational subject S, any act ϕ , any two consecutive times t_1 and t_2 , and any later time t_n ($2 < n$), S's ϕ -ing at t_n is, as of t_1 , obligatory if and only if, and because, every set of attitudes that S is, as of t_1 , permitted to have at t_2 has a 100% objective chance of resulting in S's intentionally ϕ -ing at t_n .²

² From this, we can derive the following accounts of permissibility and impermissibility. S's ϕ -ing at t_n is, as of t_1 , permissible if and only if, and because, there is some set of attitudes that S is, as of t_1 , permitted to have at t_2 that has some objective chance of resulting in S's intentionally ϕ -ing at t_n . S's ϕ -ing at t_n is, as of t_1 , impermissible if and only if,

Before I attempt to defend attitudism, I need to clarify a few things about it. First, by ‘a rational subject’, I mean any subject who is responsive to both reasons and normative requirements. Such a subject will, to the extent that she approximates ideal rationality, have the following tendencies: (1) to ϕ when she’s aware of facts that constitute decisive reasons for ϕ -ing, (2) to refrain from ϕ -ing unless she’s aware of facts that constitute sufficient reasons for ϕ -ing, and (3) to abide by normative requirements. The need for 3 in addition to both 1 and 2 arises given that there are situations in which a subject has sufficient reasons for each of two or more attitudes but is prohibited by some normative requirement from having them jointly (Broome 1999). Take, for instance, the normative requirement against intending to ϕ while believing that you will not ϕ .³ And imagine that you have both sufficient reason to intend to go out to a bar with friends and sufficient reason to intend to stay home curled up with a good book, for assume that each would be equally rewarding. But now suppose that you believe that you will not go out, as you’ve already changed into your pajamas and plopped down on the couch. It is, then, impermissible for you to intend to go out so long as you maintain this belief. Yet each attitude is individually permissible. You are permitted to intend to go out, and you are permitted to believe that you will not go out. You’re just not permitted to have both attitudes simultaneously. Thus, you can be prohibited from having a certain set of attitudes not only because you lack sufficient reason for one or more of the attitudes contained in that set, but also because the set itself violates some normative requirement. And this explains the need to talk about normative requirements in addition to whether or not there is sufficient reason to have or to lack the attitudes in question.

Second, the relevant notion of chance is the objective one—the one that doesn’t depend on the agent’s epistemic position. If, on the one hand, determinism is true, the objective chance of an event is always going to be either 0 or 1. If, on the other hand, indeterminism is true, the objective chance of an event could be anywhere from 0 to 1. And it’s

and because, there is no set of attitudes that S is, as of t_1 , permitted to have at t_2 that has any objective chance of resulting in S’s intentionally ϕ -ing at t_n .

³ See Bratman 1987 (p. 38), Harman 1976 (p. 432), Portmore 2011 (p. 174), Sobel 1994 (p. 239), Velleman 1989, and Wallace 2001 (p. 20).

the objective chance that matters, even if we're interested in the agent's subjective obligations.⁴ To see why, suppose that Kazumi, an officer in the LAPD Bomb Squad, has, as of t_1 , a subjective obligation to have the set of attitudes, A_1 , at t_2 , which includes both the desire to deactivate the bomb and the justified belief (justified by her training and experience) that cutting the red wire will deactivate it. And suppose that given some indeterminacy in the world Kazumi's having A at t_2 has a 10% (0.1) objective chance of resulting in Kazumi's intentionally cutting the green wire at t_3 . In that case, I think that we must conclude that it is subjectively permissible for Kazumi to cut the green wire at t_3 . If it's the indeterminacy in the world that leads to her intentionally cutting the green wire, then it's the world, not her, that is to blame for her cutting the green wire.

Third, the key question as far as attitudism is concerned is whether there is some set of attitudes that S is, as of t_1 , permitted to have at t_2 that has some objective chance of resulting in S 's *intentionally* ϕ -ing at t_n . Thus, it is important to note that only a set of attitudes that includes something such as an intention to ϕ could result in S 's *intentionally* ϕ -ing. This is because what distinguishes a bodily movement that constitutes an intentional action (e.g., my foot rising as a result of my intention to kick a football) from an otherwise identical bodily movement that constitutes a mere happening (e.g., my foot rising as a result of a doctor's striking my patellar ligament with her reflex hammer) is that the former, and not the latter, is appropriately (i.e., non-deviantly) caused by the relevant sort of mental state (or states), such as the intention to perform the bodily movement in question.⁵ So suppose that I'm obligated to intend to remain perfectly still, and further suppose that this intention has a 100% objective chance of causing involuntary muscle contractions in my index finger, resulting in its twitching. It doesn't follow, on attitudism, that I'm obligated

⁴ If we're interested in an account of S 's subjective obligations, then we need only insert the word 'subjective' before both 'obligatory' and 'permitted' in the above formulation of attitudism. If, instead, we're interested in an account of S 's objective obligations, then we should instead insert the word 'objective' before both 'obligatory' and 'permitted'. In either case, though, we leave the word 'objective' before 'chance'.

⁵ Besides the intention to ϕ , the relevant sorts of mental states may include a belief-desire pair (such as the desire for some end and the belief that that end would be achieved by one's ϕ -ing) and, perhaps, even just a normative belief (such as the belief that one is obligated to ϕ). I want to allow, then, that there can be various sorts of mental states that both cause bodily movements (or movements of the mind) and are such that, when they do, the movements that they effect count as intentional actions. Nevertheless, I will often just talk about the intention to ϕ , using it as a kind of placeholder for any relevant sort of mental state.

to twitch my index finger, for although the required intention has a 100% objective chance of resulting in the twitching of my index finger, it has a 0% objective chance of resulting in the *intentional* twitching of my index finger.

Fourth, in the above formulation of attitudism, I index obligations to times. That is, I speak of an agent's being obligated, *as of a certain time*, to ϕ . This is because an agent can be obligated, as of a certain time, to ϕ only if her ϕ -ing is, *as of that time*, an option for her. And since what was once an option for her may no longer be one, what was once required of her may no longer be required. To illustrate, imagine that, last week, Alejandra had the option of enlisting in either the Army or the Navy. But now that she has enlisted in the Army, she can't enlist in the Navy, for this isn't permitted. And if her options can change over time, then so too can her obligations. Suppose, for instance, that Alejandra had promised her father a month ago that she would enlist in the Navy on her birthday, which is today. A month ago, then, she had an obligation to enlist in the Navy on her birthday. But, as of yesterday, when she enlisted in the Army, she no longer has the option of enlisting in the Navy. Thus, she no longer has an obligation to do so. Of course, she may have an obligation to apologize to her father for breaking her promise. But she cannot now be obligated to do what isn't even an option for her. This means that, if we are to be careful, we must make explicit the relevant temporal indices when talking about options and obligations.

Having clarified these matters, I can now go on to defend attitudism. My strategy will be to start by defending a rather weak claim about obligations with respect to attitudes and then defend ever-stronger claims until we arrive at attitudism. The progression of claims will go as follows: (C1) we have obligations with respect to attitudes; (C2) one such obligation is the obligation to refrain from having impermissible sets of attitudes; (C3) our obligation to refrain from performing impermissible acts derives from this more fundamental obligation to refrain from having impermissible sets of attitudes; and, more specifically, (C4) our ϕ -ing at t_n is, as of t_i , obligatory if and only if, and because, every set of attitudes that we are, as of t_i , permitted to have at t_2 has a 100% objective chance of resulting in our intentionally ϕ -ing at t_n .

In sections 1–4, I defend each of these four claims in turn. In section 5, I defend two propositions about when we are permitted to have various attitudes. These propositions are then used to supplement attitudism so that it yields specific verdicts about what agents are obligated to do. And this enables me, in sections 6–8, to apply attitudism to three puzzling cases of rational choice, which in turn allows me to explore three important implications of attitudism. I then sum up my conclusions and discuss their importance in section 9.

1. Obligations with Respect to Attitudes

It is at least *prima facie* plausible to suppose that we have obligations with respect to attitudes. After all, we often claim that there are such obligations. Consider, for instance, claims i–vi with which I began this paper. Consider also that it seems quite natural to suppose that people are obligated to want what’s best for their children and to believe what they have conclusive evidence for believing. The only reason that some philosophers deny that we have such obligations is that they hold both that we do not have voluntary control over our attitudes and that we must have voluntary control over them in order to have obligations with respect to them.

For a subject to have voluntary control over whether or not she ϕ s is for her to have the same sort of control over whether or not she ϕ s that I have over whether or not I, say, intentionally touch my nose. In other words, for a subject to have voluntary control over whether or not she ϕ s is for her to have the capacity to form an effective intention to ϕ , one that would directly result in her intentionally ϕ -ing. And, clearly, we don’t generally have this sort of control over our attitudes.⁶ Whereas whether I will touch my nose just depends on whether I intend to do so, whether I will, say, form the belief that I have plenty of milk in the fridge doesn’t. Of course, I may perform an act that then results in my forming this belief. I may, for instance, perform the act of opening the fridge and looking inside, and this act may result in my forming the belief that I have plenty of milk. But it is

⁶ My claim is only that believing, desiring, and intending are not things that we generally do at will. For all that I claim, then, it may be that pretending, imagining, supposing, and deciding are things that we do at will—see Clarke 2008.

the act of opening the fridge, not the formation of the belief, that's under my voluntary control. I form the belief that I have plenty of milk involuntarily in response to my perception of a full gallon jug of milk in the fridge.

Since we don't seem to have voluntary control over our attitudes, some argue that we cannot have obligations with respect to them. The argument goes as follows:

1. For any ϕ , S's ϕ -ing is obligatory only if S can, in the relevant sense, ϕ , for 'obligated to ϕ ' implies 'can, in the relevant sense, ϕ '.
2. For any ϕ , S can, in the relevant sense, ϕ only if S has voluntary control over whether or not she ϕ s.
3. For any attitude or set of attitudes ϕ , S does not have voluntary control over whether or not she ϕ s.⁷
4. Therefore, neither individual attitudes nor sets of attitudes can be obligatory.

The problem with this argument is that premise 2 is unmotivated. Of course, it would be perfectly appropriate to infer 'S has voluntary control over her ϕ -ing' from 'S can, in the relevant sense, ϕ ' if we were to restrict the range of ' ϕ ' to actions. But the appropriateness of the inference when ' ϕ ' is restricted to actions has nothing to do with 'can' and everything to do with the nature of actions.⁸ After all, what distinguishes an action from a mere happening is, as we noted above, that only the former is the type of event that was under the subject's voluntary control. My foot rising as a result of my intention to kick a football is an action, whereas my foot rising as a result of a doctor's striking my patellar ligament with her reflex hammer is a mere happening. And this is because the former, but not the latter, was under my voluntary control—that is, the former, but not the latter, was caused by my intending to kick. So if we restrict ' ϕ ' to actions, then S will necessarily have voluntary control over her ϕ -ing simply because actions are events over which the subject has voluntary control. But this leaves it unclear as to why we should accept the inference from 'S can, in the relevant sense, ϕ ' to 'S has voluntary control over her ϕ -ing' if ' ϕ ' is not restricted to actions, which it isn't in premise 2. And premise 2 needs to be well motivated

⁷ This is controversial, but I'll grant it for the sake of argument.

⁸ This idea is borrowed from Chuard and Southwood 2009. For other similar responses to this sort of argument, see Graham 2014 (p. 400, n. 22) and Hieronymi 2008.

given that it flies in the face commonsense, for we typically hold that agents are obligated to have certain attitudes (such as beliefs) and thus can, in the relevant sense, have these attitudes despite lacking voluntary control over them.

But maybe this has been too quick. Perhaps, premise 2 could be motivated by appeal to the notion of responsibility. The thought would be that S cannot be obligated to ϕ unless ϕ is the sort of thing for which she can be appropriately held responsible and that she cannot be appropriately held responsible for having ϕ ed unless she had voluntary control over whether or not she ϕ ed. But although it is plausible to assume that S must have had some sort of control over whether or not she ϕ ed for it to be appropriate to hold her responsible for having ϕ ed, I don't see any reason for thinking that the relevant sort of control must be voluntary control. After all, we often hold people responsible for their beliefs, desires, and intentions, and this seems entirely appropriate despite the fact that these attitudes are not under their voluntary control. As Scanlon notes, "Because 'being responsible' is mainly a matter of the appropriateness of demanding reasons, it is enough that the attitude in question be...one that either directly reflects the agent's judgment or is supposed to be governed by it. For this reason, one can be responsible not only for one's actions but also for intentions, beliefs, and other attitudes" (Scanlon 1998, p. 22). Thus, the relevant sort of control seems to be *rational control*, where S has rational control over whether or not she ϕ s if and only if whether or not she ϕ s depends on whether or not she responds appropriately to the relevant reasons, assuming that she has the general capacity to do so.⁹ And, clearly, we have this sort of control over our attitudes. And since we have no reason to think that the relevant sense of 'can' necessitates voluntary control as opposed to rational control (at least in the case of attitudes), we should conclude the premise 2 is unmotivated.

What's worse is that premise 2 leads to a vicious regress. According to premise 2, S can, in the relevant sense, ϕ only if S has voluntary control over whether or not she ϕ s. But what is it for S to have voluntary control over whether or not she ϕ s? Clearly, one thing that's necessary is that she can effectively intend to ϕ such that she would ϕ if she were to

⁹ I borrow the general notion of rational control from A. M. Smith 2010, but this particular formulation of the notion is my own.

intend to ϕ . But just as clear is that this is not sufficient. For its being sufficient would imply that unconscious people can do all sorts of things that they can't do. After all, in the nearest possible world in which an unconscious, but otherwise able-bodied, patient *intends* to raise her arms (a world in which she is conscious), she raises her arms. Thus, the counterfactual "the unconscious patient would raise her arms if she were to intend to do so" is true. So if this counterfactual's being true were sufficient for such a patient to have voluntary control over whether or not she raises her arms, then we would have to accept the absurd conclusion that this patient can raise her arms despite being unconscious and unable to form the intention to do so. To avoid such absurd implications, the proponent of premise 2 must acknowledge that the truth of the counterfactual "S would ϕ if S were to intend to ϕ " is a necessary, but not a sufficient, condition for S's having voluntary control over whether or not she ϕ s and that, in addition this necessary condition, it must be that S can, in the relevant sense, form the intention to ϕ . Thus, even though the unconscious patient would raise her arms if she were to form the intention to do so, she cannot raise her arms, because she cannot, in the relevant sense, form this intention given that she is unconscious. Now, if we accept premise 2, we must hold that she can, in the relevant sense, form the intention to raise her arms only if whether or not she forms this intention is under her voluntary control. And, as we've just discovered, for the formation of the intention to raise her arms to be under her voluntary control, it must be that she can, in the relevant sense, form the second-order intention to form the intention to raise her arms. And for the formation of this second-order intention to be under her voluntary control (as it must according to premise 2), it must be that she can, in the relevant sense, form the third-order intention to form this second-order intention. And so and so forth, and, hence, an infinite regress ensues.

Given that premise 2 is unmotivated and leads to a vicious regress, we should not be persuaded by the above argument to reject the view that there are obligations with respect to attitudes. So absent any other argument against this view (and I'm at a loss as to what it might be), we should take the view at face value. And, on its face, it (*viz.*, C1) seems quite plausible.

2. An Obligation to Refrain from Forming an Impermissible Set of Attitudes

Having established at least a presumption in favor of obligations with respect to attitudes, it should be unproblematic to add that one of these obligations is the obligation, as of t_1 , to refrain from having an impermissible set of attitudes at t_2 . This holds given that, by conceptual necessity, it is obligatory to refrain from the impermissible. So if, as I've argued, there are obligations with respect to attitudes, then it follows that one of these obligations is the obligation to refrain from having an impermissible set of attitudes, which is just C2.

3. Act-Obligations Derive from Attitude-Obligations

Of course, attitudism goes beyond the claim that we have an obligation to refrain from having impermissible sets of attitudes; it makes the further claim (viz., C3) that our obligation to refrain from performing impermissible acts derives from this more fundamental obligation. There are a couple of reasons for thinking this.

First, consider that reasons-responsiveness is the ultimate source of our obligations. It is because we have the capacity to respond appropriately to reasons—a capacity that lower animals, young children, and the criminally insane typically lack—that we have obligations for which we can be appropriately held accountable. But it is our attitudes, not our actions, that are directly responsive to reasons. Indeed, our actions are reasons-responsive only insofar as the beliefs, desires, and intentions that issue in them are reasons-responsive. For whether our attitudes issue in actions depends on things that are not at all responsive to reasons, such as the proper functioning of our nervous systems. A doctor may ask a patient with a possible spinal injury to wiggle her toes, and insofar as the patient is reasons-responsive she will form the intention to wiggle her toes. But whether this intention causes her toes to wiggle depends not at all on whether she is reasons-responsive, but on whether she has the relevant sort of spinal injury. And if, as it turns out, she is unable to control the movement of her toes due to some spinal injury, we would deny that she has an obligation to wiggle her toes. At most, we would think that she has an obligation to intend to wiggle her toes. Since reasons-responsiveness is the source of our obligations and since it is the intention to ϕ , not the act of ϕ -ing, that is directly responsive to reasons, we should think that obligations with respect to ϕ -ing derive from

obligations with respect to intending to ϕ . If we are obligated to ϕ , it is only because we are obligated to intend to ϕ and the world is such that our intention will issue in our ϕ -ing.

Second, although it's possible to have an obligation to intend to ϕ without having an obligation to intentionally ϕ , it's impossible to have an obligation to intentionally ϕ without having an obligation to intend to ϕ . The former is possible because intending to ϕ can be an option even when ϕ -ing is not, as is the case for the patient with the spinal injury who could intend to wiggle her toes but not wiggle them. The latter, by contrast, is impossible, because, as a matter of conceptual necessity, one cannot intentionally ϕ without intending to ϕ .¹⁰ Thus, intentionally ϕ -ing is logically and ontologically dependent upon intending to ϕ . And this is important, because it seems that, in general, if we have obligations with respect to both X and Y , and X is logically and ontologically dependent on Y , then our obligations with respect to Y must be more fundamental than our obligations with respect to X . To illustrate, consider that knowing that p (that is, having an appropriately justified true belief that p) is logically and ontologically dependent on believing that p . After all, it is logically impossible to know that p without believing that p . And because of this, it seems that the obligation to know that p must derive from a more fundamental obligation to believe that p , where p is both true and appropriately justified. Likewise, then, we should think that the obligation to intentionally ϕ derives from a more fundamental obligation to intend to ϕ , where ϕ is an act that one would perform if one were to intend to do so. Thus, it seems that our attitude-obligations (such as our obligation to intend to ϕ) are more fundamental than our act-obligations (such as our obligation to intentionally ϕ).

For these two reasons, we should accept C_3 —that is, we should accept that our obligation to refrain from performing impermissible acts derives from our more fundamental obligation to refrain from forming impermissible sets of attitudes. And having argued for C_1 – C_3 , I've laid the necessary groundwork for defending attitudism (or C_4). That is, I've argued that we have obligations with respect to attitudes, that one of these obligations is

¹⁰ Recall that I'm using 'intend' as a placeholder for any mental state that is such that, if it appropriately causes some bodily movement or movement of the mind, that movement counts as an intentional action.

the obligation to refrain from having impermissible sets of attitudes, and that our obligation to refrain from performing impermissible acts derives from this more fundamental obligation to refrain from having impermissible sets of attitudes.

4. *In Defense of Attitudism*

Since attitudism includes a bi-conditional, defending it requires defending the following two conditionals: (LEFT-RIGHT) if S's ϕ -ing at t_n is, as of t_1 , obligatory, then every set of attitudes that S is, as of t_1 , permitted to have at t_2 has a 100% objective chance of resulting in S's intentionally ϕ -ing at t_n and (RIGHT-LEFT) if every set of attitudes that S is, as of t_1 , permitted to have at t_2 has a 100% objective chance of resulting in S's intentionally ϕ -ing at t_n , then S's ϕ -ing at t_n is, as of t_1 , obligatory. Of course, attitudism also holds that the right side of the bi-conditional has explanatory priority, but I defended the idea that the deontic statuses of acts are to be explained in terms of the deontic statuses of attitudes when I defended C₃ above. So all that is left for me to defend is both LEFT-RIGHT and RIGHT-LEFT.

According to LEFT-RIGHT, if an agent is obligated to ϕ , then there must be no chance that she would end up doing something else as a result of her having a permissible set of attitudes. Suppose for the sake of reductio, then, that there was some chance that she would end up doing something else as a result of her having a permissible set of attitudes. Suppose, for instance, she is permitted to intend to ψ and that this intention will result in her intentionally ψ -ing. How, then, could we hold her obligated to ϕ and criticize her for instead ψ -ing when she ψ -ed as a result of her permissibly intending to ψ ? The answer is that we can't, and, thus, we should accept LEFT-RIGHT.¹¹

¹¹ There may seem to be an obvious counterexample to LEFT-RIGHT. Suppose that my pushing a button would cause ten people to die and have no good effects. Clearly, then, I have an obligation to not-push. But suppose that my evidence misleadingly suggests that my pushing the button is necessary to save the ten. Consequently, I believe that I must push to save the ten. And since I desire to save the ten, I form the intention to push and so push. It may seem, then, that here is a case in which I'm obligated to not-push even though I'm led to push by the permissible set of attitudes just described. But we must be careful to distinguish my subjective obligations, which are a function of my perspective on the world, from my objective obligations, which are function of the way the world is and not of how I perceive it to be. So, in this case, I would say that, subjectively speaking, I ought to believe that I must push to save the ten, and so intend to push. But, subjectively speaking, I ought to push. So if we give LEFT-RIGHT a subjective gloss, then this is no counterexample at all. For it is only, *objectively speaking*, that I ought to not-push. But if we're going to go with an objective gloss of LEFT-RIGHT, then we should deny that you ought to believe that you must push to save the ten. Instead, you objectively ought to believe what's true, which is that your pushing will kill the ten. And if you were to have this belief,

So there is a substantial cost to denying LEFT-RIGHT: doing so requires us to accept that someone can be blameworthy for having ψ -ed even though she was not blameworthy for having the attitudes that resulted in her intentionally ψ -ing. Even so, one could deny attitudism by instead denying RIGHT-LEFT. But, as we'll see presently, this is not an attractive option either. According to RIGHT-LEFT, if having some permissible set of attitudes guarantees that one will ϕ so that the only way that one would not ϕ is by having some impermissible set of attitudes, then one is obligated to ϕ . For if the only way to see to it that one doesn't ϕ is to have some impermissible set of attitudes, then one either ϕ s or forms an impermissible set of attitudes. And since the latter is impermissible, it follows that ϕ -ing is the only permissible option. And to say that ϕ -ing is the only permissible option is just to say that ϕ -ing is obligatory. Thus, we should accept RIGHT-LEFT.

So I've argued that we should accept each of attitudism's components: LEFT-RIGHT, RIGHT-LEFT, and the explanatory priority of the right side of the bi-conditional. But another reason to accept attitudism is, as we'll see, that it has plausible implications in a variety of cases. But before we can explore its implications, we must first supplement it with some propositions about when it is permissible to have a set of attitudes, a task to which I now turn.

5. *Two Supplementary Propositions*

Recall that the view that I endorse is

Attitudism: For any rational subject S, any act ϕ , any two consecutive times t_1 and t_2 , and any later time t_n ($2 < n$), S's ϕ -ing at t_n is, as of t_1 , obligatory if and only if, and because, every set of attitudes that S is, as of t_1 , permitted to have at t_2 has a 100% objective chance of resulting in S's intentionally ϕ -ing at t_n .

Obviously, this view cannot tell us whether S is obligated to ϕ at t_n unless we have some way of determining whether or not S is permitted to have a given set of attitudes. We can, I believe, best determine this by appealing to the following account:

you would form the intention to not-push and so not-push, which is exactly what you objectively ought to do. So regardless of whether we give LEFT-RIGHT an objective gloss or a subjective gloss, this is no counterexample to it.

- ATT** For any rational subject S, any set of *attitudes* \mathbb{A} , and any two consecutive times t_1 and t_2 , S is permitted, as of t_1 , to have \mathbb{A} at t_2 if and only if
- (a) S has, as of t_1 , the option of having \mathbb{A} at t_2 ,
 - (b) S has, as of t_1 , sufficient reason to have, at t_2 , each of the attitudes in \mathbb{A} ,
and
 - (c) S's having \mathbb{A} at t_2 wouldn't violate any normative requirements.

Although ATT is fairly straightforward, it does contain three clauses that need some explaining. First, there is clause *a*, whose purpose is to exclude the possibility that a subject could be required to come to have a set of attitudes that she couldn't possibly form given that she is, say, unconscious, lacking the necessary conceptual apparatus, or subject to manipulation by a Frankfurtian demon who would prevent her from forming the relevant attitudes if she were otherwise going to form them (Frankfurt 1969). Second, there is clause *b*, which should be uncontroversial, for, by definition, it is impermissible for S to ϕ if S lacks sufficient reason to ϕ . That the reasons for S to ϕ are insufficient to make S's ϕ -ing permissible is just what it means for S to lack sufficient reason to ϕ . And, lastly, there is clause *c*, whose purpose should be apparent given the point that I made in the opening part of this paper about the possibility of situations in which a subject has sufficient reason to have each of two or more attitudes but is prohibited by some normative requirement from having them jointly.

Merely adding ATT to attitudism isn't enough to get us substantive verdicts about the permissibility of actions. For that, we would additionally need an account of when there is sufficient reason to have various types of attitudes as well as an account of the various normative requirements that govern these attitudes. And, unfortunately, doing all that is far beyond the scope of this paper. So when it comes to exploring the implications of attitudism in the following section, I'll need to tread carefully, picking only cases in which it is clear that the relevant sets of attitudes meet clauses *a–c* of ATT. Nevertheless, I will endeavor to provide an account of when there is sufficient reason to intend to perform an action. Actually, I'll provide such an account for only a narrow range of cases, cases in which the only thing that is relevant to S's decision of what to do at t_n is how much utility her various alternatives might produce. I'll call such cases *utility cases*.

The following account is, regrettably, rather complicated. So let me start off by explaining the motivation behind it. I'll then try to explain the gist of the account before proceeding to state it in full. And, after stating it in full, I will, of course, explain the details. But let's start with the motivation behind it.

In utility cases, whether S should ϕ just depends on whether S's ϕ -ing would maximize utility. But, interestingly, whether S's ϕ -ing would maximize utility often depends on what background attitudes she has. (An attitude is a background attitude with respect to S's ϕ -ing if and only if it is some attitude other than the intention of S to ϕ .) To illustrate, suppose that Professor Procrastinate has just received two emails requiring his immediate attention. One is from a journal, inviting him to write a book review.¹² The other is from his department chair, asking him to commit to a topic for the seminar that he'll be teaching next semester. If he accepts the invitation while intending to respond to his chair with a commitment to make the book the topic of his seminar, then he will read the book and write the review. However, if he accepts the invitation while intending to respond to his chair with a commitment to teach a seminar on Kant (a topic over which he tends to obsess), he will not read the book or write the review, but will instead obsess over how to interpret Kant. Thus, his accepting the invitation to write the book review will have good consequences if and only if he intends to respond to his department chair with a commitment to teach his seminar on the book.

To take one other example, suppose that Gopal has just been offered the chance to ride one of the Royal Lipizzaner Stallions. He is warned, though, that this horse can sense fear and will throw any rider who he senses to be fearful. Now, Gopal has no reason to fear the horse, but if he does, he should decline the offer, for it would be bad to get thrown. If, however, he rightly has no fear of the horse, then he should accept this rare opportunity to ride such a magnificent animal. So here too is a case where whether someone should ϕ (e.g., ride this horse) depends on what his or her background attitudes are (e.g., whether he fears this horse).

As these examples illustrate, it is important to evaluate actions in the context of the agent's background attitudes. This means that whether an agent has sufficient reason, say,

¹² This is a revised version of a case that I borrow from Jackson & Pargetter 1986.

to intend to accept an invitation to write a book review or to ride a certain horse depends on what her background attitudes are. And this need to accommodate the importance of considering S's background attitudes when assessing whether she has sufficient reason to intend at t_2 to ϕ at t_n is the main motivation for the account given below.

That's the motivation, and the gist of the account is that, in utility cases (cases in which utility is the only salient feature), S has sufficient reason to intend at t_2 to ϕ at t_n if and only if there is some permissible set of background attitudes B_1 such that the utility resulting from S's having B_1 at t_2 and ϕ -ing at t_n is at least as great as that resulting from any other combination of acts and permissible background attitudes. Thus, in the case of Gopal, he has sufficient reason to ride the horse, because there is some permissible set of background attitudes—specifically, the one in which he has no fear of the horse—such that the utility resulting from his having these background attitudes and riding the horse is at least as great as (indeed, greater than) that resulting from his having a permissible set of background attitudes and not riding the horse.

That's the gist. More formally, though, the account is as follows.

INT For any rational subject S in a utility case, any act ϕ , any two consecutive times t_1 and t_2 , and any later time t_n ($2 < n$), S has, as of t_1 , sufficient reason to *intend* at t_2 to ϕ at t_n if and only if, and because, there is some permissible set of background attitudes B_1 such that (1) S would ϕ at t_n if S were to intend at t_2 to ϕ at t_n while possessing B_1 at t_2 and (2) there is no alternative act ψ and permissible set of background attitudes B_2 —where B_2 may or may not be identical to B_1 —such that (a) S would ψ at t_n if S were to intend at t_2 to ψ at t_n while possessing B_2 at t_2 and (b) the prospect of S's ψ -ing at t_n in the context of having possessed B_2 at t_2 is better than the prospect of S's ϕ -ing at t_n in the context of having possessed B_1 at t_2 .¹³

There is much that needs to be explained here. First, we have conditions 1 and 2a, which hold S would ϕ/ψ at t_n if S were to intend at t_2 to ϕ/ψ at t_n while possessing B_1/B_2 at t_2 . This is needed in order to avoid the possibility that agents could be required to form

¹³ The prospect of S's possessing B_2 at t_2 and ψ -ing at t_n is better than the prospect of S's possessing B_1 at t_2 and ϕ -ing at t_n if and only if $\sum_i [Pr(o_i/B_2\psi) \times U(o_i)]$ is greater than $\sum_i [Pr(o_i/B_1\phi) \times U(o_i)]$, where o_i is a possible outcome resulting from S's possessing B_1 at t_2 and χ -ing at t_n , $Pr(o_i/B_1\chi)$ is the objective probability of o_i resulting from S's possessing B_1 at t_2 and χ -ing at t_n , and $U(o_i)$ is the total utility in o_i .

ineffective intentions. And, as we've seen from the case of Professor Procrastinate, whether an intention (such as the intention to write a book review) will be effective or not depends on what one's background attitudes are. If Professor Procrastinate intends to write the book review while having the background attitude of intending to commit to teaching a seminar on the book, then the intention to write the book review will be carried out. But if Professor Procrastinate intends to write the book review while having the background attitude of intending to commit to teaching a seminar on Kant, then the intention to write the book review will not be carried out.

Second, I've restricted the relevant background attitudes to permissible ones. This is because the attitudes in question are ones that the agent can acquire only involuntarily in the response to reasons and/or normative requirements. Thus, it wouldn't make sense to require agents to come to have attitudes that they weren't permitted to have. For, in that case, they wouldn't be able to acquire them involuntarily by responding appropriately to their reasons and/or normative requirements. Nor would they be able to acquire them voluntarily by intending to form them. For these attitudes, as discussed above, cannot be acquired voluntarily—at least, not typically. Thus, we should restrict the relevant background attitudes to those that are permissible. And this means, for instance, that we can require Professor Procrastinate to have the background attitude of intending to teach a seminar on the book, but not require Professor Procrastinate to have the background attitude of believing (perhaps, due to wishful thinking) that writing the book review will land him the White's Chair of Moral Philosophy at Oxford. Both of these two background attitudes may be sufficient to render his intention to write the book review effective, but we cannot require Professor Procrastinate to have impermissible background attitudes such as the belief that writing a simple book review could somehow land him such a prestigious position at Oxford.

Third, condition 2 is stated in terms of the *prospect* of S's possessing B_i at t_2 and χ -ing at t_n as opposed to the *consequences* of S's possessing B_i at t_2 and χ -ing at t_n , because there may be no determinate fact as to what the consequences would be. After all, indeterminism may be true. And I talk about the prospect of S's possessing B_i at t_2 and χ -ing at t_n as opposed to the prospect of S's possessing B_i at t_2 and *intending to* χ at t_n (the relevant con-

trast being between *performing* χ and *intending to perform* χ), because whether there is sufficient reason to intend to χ depends, not on the utility associated with S's intending to χ , but on the utility associated with S's χ -ing. The fact that I would receive a million dollars tomorrow morning if, at midnight tonight, I were to intend to drink some toxin tomorrow afternoon doesn't give me any reason to drink the toxin tomorrow afternoon. Indeed, the fact that my drinking the toxin tomorrow afternoon would have no benefit (as my receiving the million dollars depends only on what I intend to do at midnight and not on whether I drink the toxin tomorrow afternoon) and would only make me terribly ill is a reason for me *not* to drink the toxin tomorrow afternoon.¹⁴

The rest should be self-explanatory, although exactly how INT functions and why we should accept it may not be entirely clear until we have worked through a couple of examples. So, in the following three sections, I explore how the view that conjoins attitudism with both ATT and INT—which I'll call *conjunctive attitudism*—applies in three cases. These cases will illustrate some important implications of conjunctive attitudism.

6. *The Application and Implications of Conjunctive Attitudism: The Buttons*

The first case that I want to discuss is a modified version of a case that comes from Donald Regan's excellent book *Utilitarianism and Co-operation* (1980). It's an interesting case and is particularly important in motivating conjunctive attitudism, because it shows that an adequate practical theory will not be *exclusively act-orientated* but will instead be *attitude-dependent*. To say that a theory is not exclusively act-orientated is to say that it requires something more of agents than just the performance of certain voluntary acts.¹⁵ For instance, conjunctive attitudism is not exclusively act-orientated in that it requires agents to involuntarily form certain attitudes as well as to perform certain voluntary acts. And to say that a theory is attitude-dependent is to say that it takes what agents ought to do to be a function of what attitudes they ought to have. Conjunctive attitudism is attitude-

¹⁴ This case is borrowed from Kavka 1983.

¹⁵ 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 had in mind.

dependent in that it holds that what agents ought to do is a function of what they would do if they were to have the attitudes that they ought to have.

I'll call my modified version of Regan's case *The Buttons*, because it involves two individuals, Coop and Uncoop, each with a button in front of them. Depending on whether each pushes his button at t_4 , the consequences will be as depicted in Table 1. Assume that the resulting utiles will be evenly distributed over all parties, including both Coop and Uncoop.

Table 1

<i>The Buttons</i>	Uncoop: Push at t_4	Uncoop: Not-push at t_4
Coop: Push at t_4	100 utiles	0 utiles
Coop: Not-push at t_4	0 utiles	60 utiles

Coop (the cooperative one) desires that they both push at t_4 . But, unfortunately, Uncoop (the uncooperative one) lacks this desire. And Coop will push his button at t_4 if and *only if* he finds at t_3 that Uncoop desires that they both push. Let me just stipulate that they will each be able to read the other's mind at t_3 and thereby come to know what the other desires.¹⁶ So Uncoop will not push at t_4 , for he lacks the desire to push. And Coop will not push at t_4 because he will read Uncoop's mind at t_3 and find that he lacks the desire that they both push. So, neither Coop nor Uncoop is going to push at t_4 , but we can still ask of each of them: "Which act, push or not-push, is he (as of t_1) required to perform at t_4 ?"

According to what I'll call *the standard view*, the answer is surprisingly that each of them is (as of t_1) required to not-push at t_4 , for the following is

The Standard View: For any rational subject S in a utility case, any act ϕ , any time t_1 , and any later time t_n ($1 < n$), S's ϕ -ing at t_n is, as of t_1 , obligatory if and only if, and because, there is no alternative act ψ such that the prospect of ψ -ing at t_n is at least as

¹⁶ If you don't like fanciful examples involving mind readers, then just imagine that Coop and Uncoop are both poker players, that Uncoop has an involuntary tell with respect to his desires (one that he cannot fake), and that Coop is able to read Uncoop's tell.

good as the prospect of ϕ -ing at t_n .¹⁷

The standard view implies that Coop is (as of t_1) required to not-push at t_4 , because, given that Uncoop is going to not-push at t_4 , the prospect of Coop's pushing at t_4 is 0 utiles, whereas the prospect of Coop's not-pushing at t_4 is 60 utiles. (I'm assuming for the sake of simplifying the discussion that determinism is true and, thus, that all the relevant objective probabilities are either 0 or 1.) And the standard view also implies that Uncoop is (as of t_1) required to not-push at t_4 . For given that Uncoop won't desire at t_3 that they both push, it follows that Coop isn't going to push at t_4 . And given that Coop isn't going to push at t_4 , the prospect of Uncoop's pushing at t_4 is 0 utiles, whereas the prospect of Uncoop's not-pushing at t_4 is 60 utiles. So, according to the standard view, Coop and Uncoop are both required (as of t_1) to not-push at t_4 .

Note that it is because the standard view is exclusively act-orientated that it cannot require Uncoop to "do" all that is necessary to ensure the optimal outcome. To ensure the optimal outcome, Uncoop must do more than just perform the voluntary act of pushing. Uncoop must additionally form the desire that they both push, and forming this desire is not a voluntary act. Indeed, the only way that Uncoop can come to desire that they both push is by recognizing and responding appropriately to the decisive reason he has to form this desire—that reason being that their both pushing would produce the optimal outcome. Thus, being exclusively act-oriented, the standard view is unable to require Uncoop to involuntarily form the desire that they both push. Of course, there may seem to be an easy and obvious fix: simply conjoin the standard view with ATT. Call this the conjunctive version of the standard view. But this won't do, for although the conjunctive version of the standard view requires Uncoop to desire that they both push, it doesn't require Uncoop to push. Given that Uncoop is not going to form the requisite desire, the conjunctive version of the standard view implies that Uncoop should not-push. But what we need is a theory that requires that Uncoop both desires at t_3 that they both push and pushes at t_4 . For only

¹⁷ The prospect of ψ -ing at t_n is at least as good as the prospect of ϕ -ing at t_n if and only if $\sum_i [Pr(o_i/\psi) \times U(o_i)]$ is greater or equal to $\sum_i [Pr(o_i/\phi) \times U(o_i)]$, where o_i is a possible outcome resulting from S's χ -ing at t_n , $Pr(o_i/\chi)$ is the objective probability of o_i resulting from S's χ -ing at t_n , and $U(o_i)$ is the total utility in o_i .

a theory that requires both of Uncoop will require Uncoop to “do” all that is necessary to secure the optimal outcome.

Nevertheless, as I will show presently, even this is not enough. That is, it isn’t enough for a theory to require Uncoop both to desire at t_3 that they both push and to push at t_4 . An adequate practical theory must be attitude-dependent. That is, it must hold that what an agent ought to do depends on what attitudes she ought to have. To see why, consider the following view:

The Cooperative View: For any rational subject S in a utility case, any act ϕ , any time t_i , and any later time t_n ($1 < n$), S’s ϕ -ing at t_n is, as of t_i , obligatory if and only if, and because, S’s ϕ -ing at t_n is part of the best cooperative scheme.

This view implies that Uncoop should push, for Uncoop’s pushing is part of the best cooperative scheme: the one in which Coop and Uncoop both push. So if we were conjoin the cooperative view with ATT (call this the conjunctive version of the cooperative view), we would end up with a theory that requires Uncoop both to desire at t_3 that they both push and to push at t_4 . Thus, we would end up with a theory that requires Uncoop to “do” all that is necessary to secure the optimal outcome in *The Buttons*. However, this is not enough. For such a theory would have counterintuitive implications in the following variant on *The Buttons*. Call it *The Frankfurtian Buttons*. In this case, everything is as it was in *The Buttons* except that there is a Frankfurtian demon standing by to prevent Uncoop from forming the desire that they both push if he were otherwise going to form this desire. Thus, in this case, condition *a* of ATT (which holds that Uncoop must have the option of forming the desire in order to be required to form it) is not met. So, in *The Frankfurtian Buttons*, Uncoop is not required to form the desire that they both push. Yet the conjunctive version of the cooperative view still requires Uncoop to push, because his pushing is still part of the best cooperative scheme: the one in which he and Coop both push. But, given that Uncoop’s pushing without desiring that they both push would have terrible consequences, we should think that Uncoop is required to push only if Uncoop is also required to form the desire that they both push. The problem, then, with the conjunctive version of the cooperative view is that it requires Uncoop to push even if he is not re-

quired to desire that they both push. What we need instead is an attitude-dependent view, a view that makes what Uncoop is required to do dependent on what attitudes he is required to have.

So we've looked at four alternatives to conjunctive attitudism: (1) the standard view, (2) the cooperative view, (3) the conjunctive version of the standard view, and (4) the conjunctive version of the cooperative view. The problem with first two is that they are exclusively act-orientated and so can't require Uncoop to desire that they both push. The problem with the conjunctive version of the standard view is that although it requires Uncoop to desire that they both push, it fails to require Uncoop to push. And the problem with the conjunctive version of the cooperative view is that it is not attitude-dependent and, thus, holds that Uncoop should push even if he is not required to desire that they both push.

Conjunctive attitudism is superior to these four alternatives. First, not being exclusively act-orientated, it requires Uncoop, in *The Buttons*, not only to push at t_4 but also to desire at t_3 that they both push. Second, being attitude-dependent, it requires Uncoop, in *The Frankfurtian Buttons*, to not-push at t_4 given both that his desiring at t_3 that they both push is not an option and that his not-pushing at t_4 is the best that he can do if he lacks this desire. In both cases, then, conjunctive attitudism requires Uncoop to "do" all that is necessary to bring about the best outcome that he can. It's just that whereas, in *The Buttons*, that's the outcome in which they both push, in *The Frankfurtian Buttons*, that's the outcome in which they both not-push.

Just so it's clear, let me spell out how conjunctive attitudism gets these results. Let's start with *The Buttons*. Given its commitment to ATT, conjunctive attitudism implies that Uncoop is, as of t_1 , required to desire at t_3 that they both push. In this case, forming this desire is an option for Uncoop, and Uncoop has decisive reason to do so given that the object of this desire (i.e., their both pushing) has the best possible consequences. Furthermore, given conjunctive attitudism's commitment to INT, it implies that Uncoop is, as of t_1 , required to intend at t_3 to push. For, according to INT, we are to determine whether Uncoop is, as of t_1 , required to intend at t_3 to push or not-push by comparing the prospect of Uncoop's both having the requisite desire at t_3 and pushing at t_4 to the prospect of Uncoop's both having the requisite desire at t_3 and not-pushing at t_4 . And when we compare

these two, we find that the former has a value of 100 utiles, whereas the latter has a value of 0 utiles. After all, if Uncoop desires at t_3 that they both push and pushes his button at t_4 , then Coop will also push at t_4 . And if they both push, 100 utiles will result. But if Uncoop desires at t_3 that they both push but *not-pushes* at t_4 (perhaps, despite desiring that they both push, Uncoop not-pushes because he is ordered to by his commanding officer), 0 utiles will result. For if Uncoop desires at t_3 that they both push, then Coop will push at t_4 . And if Coop pushes and Uncoop doesn't, then 0 utiles will result. So conjunctive attitudism entails that Uncoop is, as of t_3 , required not only to desire at t_3 that they both push, but also to intend at t_3 to push at t_4 . And this means that conjunctive attitudism entails that Uncoop is, as of t_3 , required to push at t_4 . For every set of attitudes that Uncoop is, as of t_3 , permitted to have at t_3 has a 100% objective chance of leading to S's intentionally pushing at t_4 . For, as we've seen, all the sets of attitudes that Uncoop is, as of t_3 , permitted to have at t_3 include both the desire that they both push and the intention to push at t_4 . And I'm just assuming that the world is such that if, at t_3 , Uncoop both desires that they both push and intends to push at t_4 , this will result in his intentionally pushing at t_4 . Thus, conjunctive attitudism implies that Uncoop is, as of t_3 , required to push at t_4 .

Now, let's consider *The Frankfurtian Buttons*. In this case, ATT implies that Uncoop is *not* required to desire that they both push. Given the inclusion of clause *a*, ATT holds that an agent can be required to form a desire only if her forming the desire is an option for her. And, in *The Frankfurtian Buttons*, Uncoop's forming the desire that they both push is not an option for him. And this means that, when we apply INT, we must compare the prospect of Uncoop's both lacking the desire at t_3 that they both push and pushing at t_4 to the prospect of Uncoop's both lacking the desire at t_3 that they both push and not-pushing at t_4 . And when we compare these two, we find that the former has a value of 0 utiles, whereas the latter has a value of 60 utiles. After all, if Uncoop lacks the desire at t_3 that they both push and pushes his button at t_4 , then Coop will not-push and he will push, resulting in 0 utiles. But if Uncoop lacks the desire at t_3 that they both push but *not-pushes* at t_4 , then 60 utiles will result. So conjunctive attitudism entails that Uncoop is, as of t_3 , required to intend at t_3 to not-push at t_4 . And the world is, I'll assume, such that if, at t_3 , Uncoop both lacks the desire that they both push and intends at t_3 to not-push at t_4 , then this

will result in his not-pushing at t_4 . So conjunctive attitudism implies that, in *The Frankfurtian Buttons*, Uncoop is, as of t_3 , required to not-push at t_4 .

It seems to me, then, that conjunctive attitudism gets the intuitively right results in both *The Buttons* and *The Frankfurtian Buttons*. Yet some may question the plausibility of conjunctive attitudism's implications in *The Buttons* once they realize that it implies not only that Uncoop is, as of t_3 , required to push at t_4 but also that Coop is, as of t_3 , required to not-push at t_4 . For Coop should intend to not-push given that Uncoop is going to not-push, and I'm assuming that the world is such that if Coop intends to not-push he will not-push. But it may seem absurd for conjunctive attitudism to hold both that Uncoop should push and that Coop should not-push given that the worse possible outcome (the one with 0 utiles) will obtain if Uncoop pushes and Coop not-pushes. Contrary to how things may seem, though, this is not absurd, for Uncoop and Coop are in very different situations. Uncoop is in a position to affect what Coop does, whereas Coop is not in a position to affect what Uncoop does. If Uncoop desires at t_3 that they both push, then Coop will push. And this is why Uncoop is required both to have this desire and to push. However, Uncoop is going to not-push regardless what Coop does or desires. And this is why Coop must just take for granted that Uncoop will not-push and make the best of a bad situation by not-pushing as well.¹⁸

Of course, some people could, despite all these arguments, maintain that Uncoop should not-push given that Coop is going to not-push, and they may even acknowledge that Coop's not-pushing is contingent upon Uncoop's failing to fulfill his obligation to desire that they both push.¹⁹ But if they do, they should also acknowledge that their position implies that Uncoop faces a normative dilemma for which there is no way out. For if Un-

¹⁸ If you still have any doubts, then just imagine the following variation of *The Buttons: Buttons 2*. This case is exactly like *The Buttons* except that in this case whether Coop will push depends, not on whether Uncoop desires that they both push, but on whether Uncoop tells Coop that he's willing to push. And in *Buttons 2* Uncoop is not going to tell Coop that he's willing to push. In this case, it seems that Uncoop should push for he should both tell Coop at t_3 that he's willing to push and push at t_4 . But Coop should not push, for Uncoop is not going to push and nothing he can do will change that.

¹⁹ Also, those who hold that Uncoop should not-push will have to deny the validity of deontic detachment, which holds that, from $O(p)$ and $O(p \rightarrow q)$, $O(q)$ follows, where ' $O(x)$ ' stands for 'it ought to be that x ' (Greenspan 1975). After all, it ought to be that Uncoop desires that they both push, and it ought to be that, if Uncoop desires that they both push, he pushes. So, from deontic detachment, it follows that it ought to be that Uncoop pushes.

coop has the set of the attitudes that he ought to have (the set that includes both a desire that they both push and an intention to push), he will push. And yet these same people maintain that Uncoop should not-push. To my mind, this is too great a cost. We should not accept that normative dilemmas (for which there is no way out) can arise without their having been some past mistake—that is, some previous normative failure. And, in this case, there needn't have been any past mistake. For we can just stipulate that Coop and Uncoop arose *ex nihilo* with all and only the attitudes and dispositions that they are stipulated to have. Thus, it needn't be that Uncoop made any previous mistake that resulted in his lacking the desire that they both push. Still, it would be a future mistake for Uncoop to fail to form the desire that they both push in the face of his current awareness of decisive reasons for forming this desire.

To sum up, I've argued that we should accept both that, in *The Buttons*, Uncoop is, as of t_1 , required to push at t_4 and that, in *The Frankfurtian Buttons*, Uncoop is, as of t_1 , required to not-push at t_4 . The only kind of practical theory that can get these two intuitive verdicts is one that, like conjunctive attitudism, is not exclusively act-orientated but is instead attitude-dependent. This, I believe, constitutes good evidence in favor of conjunctive attitudism and against views like the standard view and the cooperative view. In the next two sections, I'll buttress the case in favor of conjunctive attitudism by showing that it gets intuitive verdicts in yet more cases.

7. *The Application and Implications of Conjunctive Attitudism: The Cookies*

Besides implying that an adequate practical theory must not be exclusively act-orientated and should instead be attitude dependent, attitudism implies that an agent can be obligated to perform only those acts that she has some chance of performing given a permissible set of attitudes. And this, in turn, has important implications for a key debate among practical theorists. The debate concerns following sort of case.

The Cookies: If I were to eat a cookie now, I would then eat another, and another, and so on until I've finished the whole bag, ending up sick to my stomach. But let's also suppose that if, after eating just one cookie, I decided to stop and put the bag of cookies away, I would then stop after having eaten just one cookie, which would be best.

Second best would be my eating no cookies. (And assume that if I were refrain from eating a cookie now, I would end up eating no cookies.) Third best would be my eating two cookies. Fourth best would be my eating three cookies, and so on and so forth. Worst of all would be my eating the whole bag. The problem is that although I would stop eating after having eaten just one cookie if I were at that point to decide to stop eating and put the bag away, I am in fact going to decide to continue eating after tasting how delicious they are, and this is true no matter what I intend, want, think, or feel now (before eating a cookie and finding out how delicious they are). And this is unfortunate, because this decision will lead to my eating the whole bag, which is the worst possible outcome.

On the one hand, some (e.g., Goldman 1978 and Portmore 2011) argue that my eating just one cookie is not an option for me at present given that no matter what I want, intend, think, or feel now I will not eat just one cookie. These people hold that I should refrain from eating any cookies, as this is better than my only other option, which, on this view, is my eating the whole bag. On the other hand, others (e.g., Feldman 1986 and Zimmerman 1996) argue that because I would stop after having eaten just one cookie if I were then to decide to stop and put the bag away, I currently have the option of eating just one cookie. These people, known as possibilists, hold that I should eat just one cookie and then put the bag away, for this is the best of what they take my options to be. But if attitudism is correct, then possibilists have a mistaken account of what an agent's options and obligations are. I can't now be obligated to eat just one cookie if there is no permissible set of attitudes such that, if I were to have those attitudes, I would then eat just one cookie. So if attitudism is true, then possibilism is false.

Interestingly, if conjunctive attitudism is true, then the leading contender to possibilism—namely, actualism—is also false. Actualism is just the view that I've been calling the standard view. And, on the standard view, you should eat a cookie and then put the bag away. For the prospect of your doing this is better than that of either eating no cookies or eating more than one cookie. Nevertheless, the standard view also implies that you should not eat a cookie. For the prospect of your eating a cookie (i.e., the outcome in which you end up eating the whole bag) is worse than the prospect of your not eating a cookie (i.e., the outcome in which you end up having eaten no cookies). So the standard view leaves

you with no way out. If you eat a cookie, you will have violated the requirement not to eat a cookie. But if you don't eat a cookie, you will have violated the requirement to eat a cookie and then put the bag away. Conjunctive attitudism, by contrast, avoids such incompatible obligations, for although it agrees with the standard view that you should not eat a cookie, it denies that you should eat a cookie and then put the bag away, for, on conjunctive attitudism, this isn't even an option. After all, it's stipulated that no matter what your attitudes are now, you will not eat a cookie and then put the bag away.

8. *The Application and Implications of Conjunctive Attitudism: The Pies*

Limiting our options in this way also helps us deal with the following sort of case.

The Pies: You start off at t_0 with an apple pie (AP) and \$1. At t_1 , I offer you a blueberry pie (BP) in exchange for your apple pie plus 10¢. If you take this deal (call it Deal₁) at t_2 , I'll next offer you at t_3 a cherry pie (CP) in exchange for your blueberry pie plus 10¢. If you take this deal (call it Deal₂) at t_4 , I'll offer at t_5 to give back the apple pie with which you started in exchange for your cherry pie plus 10¢. And if you take this deal (call it Deal₃) at t_6 , you'll end up at t_7 with the same apple pie with which you started, only 30¢ poorer. If you don't take a given deal, you won't be offered any subsequent deals.

Unfortunately, you have the following cyclical and, thus, intransitive preferences:

$$\{AP, \$1\} < \{BP, 90¢\} < \{CP, 80¢\} < \{AP, 70¢\} < \{AP, \$1\} < \{CP, 80¢\}$$

Here, 'A < B' stands for 'you prefer B to A' and '{x, y}' stands for 'the state of affairs in which you have both x and y'.

Given these preferences, you are inclined at t_1 to take Deal₁ at t_2 . But, if you take Deal₁ at t_2 , you will subsequently take both Deal₂ at t_4 and Deal₃ at t_6 . Even if you were to intend at t_1 to take only the first two deals, you would at t_5 change your mind and decide to take Deal₃ at t_6 , ending up, then, with the same apple pie with which you started, only 30¢ poorer. Indeed, the only way to ensure that you won't later change your mind is to impermissibly form at t_1 the intention to kill whoever has the apple pie at t_7 , whether that be me or you. So if you take Deal₁ at t_2 (assuming that you don't intend to kill whoever has the apple pie at t_7), you'll end up with {AP, 70¢}. And if you refuse Deal₁ at t_2 , you'll end up with {AP, \$1}.²⁰

²⁰ This is adapted from a similar case found in Hedden 2013.

Many would argue that such cyclical and intransitive preferences are irrational. Some argue this, because cyclical preferences can, as this example illustrates, turn you into a money pump. Others argue this, because they think that our preferences should track the betterness relation, which they claim is transitive. If cyclical preferences are indeed irrational, then it's fortunate that conjunctive attitudism is not exclusively act-orientated and can, therefore, require you to revise your preferences. Suppose, then, that if you were to revise your preferences so as to make them rational, you would come to have the following transitive preferences:

$$\{AP, 70¢\} < \{AP, \$1\} < \{BP, 90¢\} < \{CP, 80¢\} \not< \{AP, 70¢\}$$

In this case, you don't have to worry about ending up with $\{AP, 70¢\}$. Since you prefer $\{CP, 80¢\}$ to $\{AP, 70¢\}$, you won't be tempted to take Deal₃. So, in this case, if you take Deal₁ at t_2 , you'll end up with $\{CP, 80¢\}$, which is your most preferred state of affairs. And if you refuse Deal₁ at t_2 , you'll end up with $\{AP, \$1\}$, which you disprefer to $\{CP, 80¢\}$. Thus, conjunctive attitudism implies that you should, in this case, revise your preferences and intend to take Deal₁ at t_2 .²¹ And since, as we'll assume, you would take Deal₁ at t_2 if this were your intention, then it follows, on conjunctive attitudism, that you should take Deal₁ at t_2 .

But what if your cyclical preferences are rational?²² In that case, conjunctive attitudism holds that you should refuse Deal₁ at t_2 . For if your cyclical preferences are rational, then we can't require you to revise them. And given such cyclical preferences, you will be sorely tempted to take each subsequent deal if you take the first one.²³ Indeed, it is stipu-

²¹ I'm assuming that this is a utility case and that your utility function is just your preference function.

²² See, for instance, Quinn 1990. And see both Andreou 2007 and Hedden 2013 for why we shouldn't think that the fact that cyclical preferences can turn us into money pumps is a good reason to think that cyclical preferences are irrational. Hedden thinks, though, that there may be better reasons for thinking that cyclical preferences are irrational. For instance, he mentions (p. 26, n. 43) that one might appeal to the transitivity of betterness and argue for a link between rational preferences and betterness. But, of course, we could, as Temkin and Rachels do, argue that betterness is not transitive and then cite the link between rational preferences and betterness as a reason to think that rational preferences needn't be transitive. See, for instance, Rachels 1998 and Temkin 1996.

²³ If you were instead a very resolute sort of person who could resist future temptation and so effectively intend at t_0 to end up at any one of the four possible states (that is, $\{AP, \$1\}$, $\{BP, 90¢\}$, $\{CP, 80¢\}$, and $\{AP, 70¢\}$), then conjunctive attitudism would imply that you face a rational dilemma, where no matter which state you choose to end up in, you will

lated that if you were to take Deal₁ at t_2 , you would then take the next two deals and end up with {AP, 70¢}, whereas, if you were to refuse Deal₁ at t_2 , you would end up with {AP, \$1}, which you prefer to {AP, 70¢}. Given that the prospect of your refusing Deal₁ at t_2 (viz., {AP, \$1}) is better than the prospect of your taking Deal₁ at t_2 (viz., {AP, 70¢}), conjunctive attitudism implies that you are required to intend to refuse Deal₁ at t_2 . And given that you will, or so I'll assume, refuse Deal₁ at t_2 if that's what you intend, it follows that, on conjunctive attitudism, you are required to refuse Deal₁ at t_2 .

You may wonder, though, why intending to refuse Deal₁ at t_2 wouldn't be prohibited given that there is an alternative—viz., the series of acts consisting in your taking both Deal₁ at t_2 and Deal₂ at t_4 but refusing Deal₃ at t_6 —whose prospect (i.e., {CP, 80¢}) is preferred to the prospect of your refusing Deal₁ at t_2 (i.e., {AP, \$1}). The answer lies with the fact that on conjunctive attitudism you can be required to perform only those acts that have some chance of being performed if you were to have some permissible set of attitudes. And, as stipulated in this example, there is no chance that you will take only the first two deals if you permissibly refrain from intending to kill whoever has the apple pie at t_7 .

So conjunctive attitudism is an attractive position with respect to *The Pies*. It's neutral on whether your cyclical preferences are irrational. But if they are irrational, then, unlike views that are exclusively act-orientated, it requires you both to change your preferences and to perform the act that you would perform were you to have the attitudes that you ought to have—viz., the act of taking Deal₁ at t_2 . And if your cyclical preferences are rational, then it holds that you are, as of t_0 , required not only (1) to refrain from forming at t_1 the intention to kill whoever has the apple pie at t_7 , but also (2) to refuse Deal₁ at t_2 . Thus, being attitude-dependent, conjunctive attitudism plausibly implies that whether you should take or refuse Deal₁ at t_2 depends on whether or not you should revise your preferences. If you should revise your preferences such that you won't be tempted at t_5 to take

have chosen wrongly, because there was an alternative state that you could have ended up in that is preferred to the one that you chose. This, I think, is as it should be. This case would then be like the case in which God will give you negative utility if you pick no positive integer and will give you a corresponding amount of positive utility if you pick some positive integer. No matter which positive integer, n , you pick, you will have picked wrongly, because it would have been better had you picked $n+1$.

Deal₃ at t_6 after you've taken both Deal₁ at t_2 and Deal₂ at t_4 , then you should, as of t_0 , take Deal₁ at t_2 . But if you are permitted to keep your cyclical preferences, then you should, as of t_0 , refuse Deal₁ at t_2 . For, given those preferences, taking Deal₁ at t_2 would just result in your ending up with the same apple pie with which you started, only 30¢ poorer.

9. Conclusion

I've argued for a radically different sort of practical theory, a theory that is not exclusively act-orientated but is instead attitude-dependent. According to such a view, we cannot understand what our practical obligations are merely by considering which acts we could perform and what features they have. We must also consider which attitudes we should have and which actions we would perform if we were to have these attitudes. This means that most typical practical theories are mistaken. Theories such as rational egoism, Rossian deontology, act-consequentialism, and rule-consequentialism are mistaken because they are exclusively act-orientated and, thus, unable to require Uncoop to have the attitudes that he is required to have—the attitudes that he must have in order to secure the optimal outcome. And theories such as global consequentialism and virtue ethics are mistaken, for even though they are not exclusively act-orientated, they fail to be attitude-dependent and thus fail to make what Uncoop should do dependent on what attitudes he ought to have. So one important implication of this paper is that most, if not all, of the theories currently on offer are inadequate and that we should, therefore, turn our attention to developing a theory such as conjunctive attitudism that meets both the constraint of not being exclusively act-orientated and the constraint of being attitude dependent.

The only two theories that I know of that are not ruled out by these two constraints are Regan's (1980) cooperative utilitarianism and Adams's (1976) conscience utilitarianism. Cooperative utilitarianism is untenable for reasons that I spell out in Portmore 2014. And conscience utilitarianism is, I believe, also untenable. Conscience utilitarianism holds both that (1) "we have a moral duty to do an act, if and only if it would be demanded of us by the most useful kind of conscience we could have" (p. 479) and that (2) we have a moral duty to have the most useful kind of conscience we could have. Note that, by 'most useful', Adams means 'most productive of utility'. One reason to think that conscience utili-

tarianism is untenable is that it morally requires agents to have attitudes that they are not even rationally permitted to have. For instance, if the belief that God punishes immoral agents with eternal damnation would be part of the most useful kind of conscience that I could have (because, say, it would motivate me to act morally), then conscience utilitarianism requires me to have this belief even if it's false and even if I have no evidence for it. To take just one other example, if a surgeon's non-instrumental desire for the death of her patient would be part of the most useful kind of conscience that she could have (because, say, it would result in her killing her patient and thereby making her donor organs available for saving the greater number), then conscience utilitarianism requires her to non-instrumentally desire the death of her patient even though the patient's death is only instrumentally valuable.

In any case, I've shown more than just that conjunctive attitudism can meet the above two constraints. I've also shown that it gets intuitively plausible verdicts in a range of otherwise puzzling cases, such as *The Buttons*, *The Cookies*, and *The Pies*. Lastly, I've shown that conjunctive attitudism has important implications with respect to what counts as an option for an agent and, consequently, offers a way of resolving the debate between actualists and possibilists.²⁴

Arizona State University
douglas.portmore@asu.edu

References

Adams, R. M. (1976). "Motive Utilitarianism." *Journal of Philosophy* 73: 467–481.

²⁴ For helpful comments and discussions, I thank audiences at Arizona State University, Colorado University, Boulder, and the 2014 Pacific Division Meeting of the American Philosophical Association. I especially thank Chrisoula Andreou, Cheshire Calhoun, David Boonin, Eric Chwang, Peter A. Graham, Johan E. Gustafsson, Chris Heathwood, Brian Hedden, Adam Hosein, Gene Mills, G. Shyam Nair, Howard Nye, Graham Oddie, Andrew Reisner, Melinda Roberts, Larry Temkin, Travis Timmerman, Michael Tooley, and Ralph Wedgwood.

- Andreou, C. (2007). "There Are Preferences and Then There Are Preferences." In B. Montero and M. D. White (eds.), *Economics and the Mind*, pp. 247-274, New York: Routledge.
- Bratman, M. E. (1987). *Intention, Plans, and Practical Reason*. Cambridge, MA: Harvard University Press.
- Broome, J. (1999). "Normative Requirements." *Noûs* 12: 398–419.
- Chuard, P., and N. Southwood (2009). "Epistemic Norms without Voluntary Control." *Noûs* 43: 599–632.
- Clarke, R. (2008). "Autonomous Reasons for Intending." *Australasian Journal of Philosophy* 86: 191–212.
- Feldman, F. (1986). *Doing the Best We Can: An Essay in Informal Deontic Logic*. Dordrecht: D. Reidel Publishing Company.
- Frankfurt, H. (1969). "Alternate Possibilities and Moral Responsibility." *Journal of Philosophy* 66: 829–839.
- Graham, P. A. (2014). "A Sketch of a Theory of Moral Blameworthiness." *Philosophy and Phenomenological Research* 88: 388–409.
- Greenspan, P. S. (1975). "Conditional Oughts and Hypothetical Imperatives." *Journal of Philosophy* 72: 259–276.
- 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.
- Harman, G. (1976). "Practical Reasoning." *Review of Metaphysics* 29: 431–463.
- Hedden, B. (2013). "Options and Diachronic Tragedy." *Philosophy and Phenomenological Research*. doi: 10.1111/phpr.12048.
- Hieronymi, P. (2008). "Responsibility for Believing." *Synthese* 161: 357–373.

- Jackson, F. and R. Pargetter (1986). "Oughts, Actions, and Actualism." *Philosophical Review* 95: 233–255.
- Kavka, G. (1983). "The Toxin Puzzle." *Analysis* 43: 33–36.
- Parfit, D. (1984). *Reasons and Persons*. Oxford: Oxford University Press.
- Parfit, D. (2011). *On What Matters. Vol. 2*. Oxford: Oxford University Press.
- Portmore, D. W. (2011). *Commonsense Consequentialism: Wherein morality meets rationality*. New York: Oxford University Press.
- Portmore, D. W. (2014). "Consequentialism and Coordination: How Consequentialism Has an Attitude Problem." Unpublished manuscript.
- Quinn, W. S. (1990). "The Puzzle of the Self-Torturer." *Philosophical Studies* 59: 70–90.
- Rabinowicz, W., and T. Rønnow-Rasmussen (2004). "The Strike of the Demon: On Fitting Pro-Attitudes and Value." *Ethics* 114: 391–423.
- Rachels, S. (1998). "Counterexamples to the Transitivity of Better Than." *Australasian Journal of Philosophy* 76: 71–83.
- Regan, D. (1980). *Utilitarianism and Co-operation*. New York: Oxford University Press.
- Scanlon, T. M. (1998). *What We Owe to Each Other*. Cambridge, Mass.: Belknap Press.
- Smith, A. M. (2010). "Attitudes and Control." Unpublished manuscript dated April 12, 2010.
- Sobel, J. H. (1994). *Taking Chances: Essays on Rational Choice*. Cambridge: Cambridge University Press.
- Temkin, L. S., (1996). "A Continuum Argument for Intransitivity." *Philosophy & Public Affairs* 25: 175–210.
- Velleman, J. D. (1989). *Practical Reflection*. Princeton: Princeton University Press.
- Wallace, R. J. (2001). "Normativity, Commitment, and Instrumental Reason." *Philosophers' Imprint* 1.

Zimmerman, M. J. (1996). *The Concept of Moral Obligation*. Cambridge: Cambridge University Press.