

# Hypothetical Oughts

Jim Pryor

20 Feb 2012

— with important long-ago help from Christian, though we'll have to see how much of this he now agrees with.

[I'm trying to reorganize the ways I discuss this material, so this will be a very unpolished presentation. I know many of you have also thought about these issues yourselves, so I hope what I'm giving you is in good enough shape for a useful discussion. And again, thanks for tolerating my hardly tolerable tardiness getting the paper to you.]

## Section 1

Our topic is: what is, and what isn't, the normative bearing of facts peculiar to the imperfectly rational. By that I don't just mean agents with mortal limits and handicaps. I mean agents who are in the grip of attitudes that are contrary to justification.

For example, suppose Adam plans to assassinate his dean. In fact, he's not justified in doing so, but that's what he plans. Yet he procrastinates about buying the poison he needs. That's not his better nature or better sense coming through; it's just his practical ineptitude. It seems, then, that although setting down this road was one practical mistake, he's compounding his failures by not taking the means he recognizes he needs to his end. Once everything is added up, perhaps he may be better off this way. Still, his practical ineptitude does seem to constitute an additional normative defect, beyond his unjustified adoption of this project in the first place. (Imagine his mother complaining, "you can't even kill someone right.")

When we recognize some means to be necessary to an end, intuitively that should establish some normative connection between them. But what exactly does that connection look like? If we say you're justified in taking means necessary to ends you just happen to have, we'd then be saying Adam is already justified in buying the poison. But that doesn't seem right. Perhaps he *would* be justified if he were justified in his assassination project. But he's not. So it's not clear he's now justified in buying the poison. That stuff is dangerous to have around the

house. Does he really have anything to counterbalance that? Or might he just think he does.

Perhaps we should say, instead, that you are only justified in taking means necessary to ends you *justifiably have*. Since Adam's end is not justified, he's not justified in taking the means either. That sounds better. But then, what underwrites our sense that he's exhibiting any normative defect in this case in not pursuing the means?

So the straightforwardly subjectivist understanding of the means-end connection seems to make buying poison too commendable for Adam. But on the other side, the straightforwardly objectivist understanding seems to let him off the hook too easily. We should want to do more justice to our impression that he's exhibiting a further defect, in not pursuing the means—without yet saying, outright, that pursuing the means is what he altogether ought to be doing.

Those are my guiding thoughts. Adam's end, though unjustified, still imposes *some kind* of normative pressure on him, and that his ignoring that pressure constitutes some kind of further normative failing. Yet this pressure seems somehow different from his having reasons to pursue the means. That doesn't seem the right way to accommodate our judgments about the case.

Of course there are lots of things that philosophers can, and do, say at this point. One idea is that Adam does after all have *some* reason to buy poison, more reason than he'd have if he didn't plan to kill the dean. It's just that he may also have stronger reasons not to buy it. We'll discuss this thought later. I'll just remark now that it doesn't seem to convincingly match the way the case strikes me and others. You can imagine Adam's commitment to the assassination project to be deeply entrenched, and his need for the poison, given that project, to be absolute. All the more puzzling why he doesn't get out to the druggist's. On the other hand, you can imagine his reasons against having the poison to be relatively mild. Now, if his merely having the end he does gave him some reason to pursue the means, surely it might in such a case tip the scales, and now make it all-things-considered better for him to buy the poison. But on the face of it, that doesn't seem to be the dynamic we think is in place. Did he have reason to buy the poison, but no reason to use it? Or did he have reason to use it, too? What reason? Didn't we start out saying his end was in fact unjustified. Something goes wrong when he doesn't buy the poison. But I don't think his normative failing here is that of acting against the balance of his reasons. The balance of his reasons may recommend keeping away from the druggist's.

Another thing people say about Adam is that his defect is one of character. He may not be doing anything wrong right now, but he's manifesting some habits, and perhaps even reinforcing them, that will serve him badly in the cases where he does have reasons to pursue his ends. Sure, that's true. But it doesn't seem to me to exhaust the sense that he's right now, occurrently, actively doing something wrong. My sense is that, there's some kind of normative pressure

here that he's right now flouting. He's not just showing off bad habits that may make him do wrong elsewhere.

I already suggested, but let me emphasize again that agents like Adam need not be *overall worse off* by making such mistakes. The best thing for Adam to do would be to realize his end is unjustified, and to abandon it. Perhaps second best would be for him to abandon the end, in a way that's less meritorious. Perhaps out of laziness. But let's suppose that abandoning the end is off the table. What then is *next* best? Is it obvious that next best is for him to carry his plan out with alacrity and efficiency? Is it better for him to be a practically consistent murderer than a harmless inefficacious guy with a grudge? That's far from obvious. I think he is displaying some normative defect, but in the circumstances, not displaying it may well amount to the worse sin.

My guiding thought here is just that Adam is now displaying some defect, that I hope to diagnose and better understand. I'm not thinking that the next best thing, if you happen to have ends you shouldn't, is always to consistently carry them out.

Some more examples. Benny intends to kill not the dean but instead her dog. And suppose this really is the right thing to do. That dog is a real hazard to the community, it needs to be done, and he's the only one willing to do it. Benny decides that the best strategy is to approach the dog slowly with a stick. In fact this is a stupid strategy, and Benny ought to know so. He really hasn't thought his options through carefully enough. But unjustified though he is, that's what he believes. He believes he ought to approach the dog slowly with a stick. Yet he doesn't. Maybe he's a bit too nervous or excited. So instead he runs headlong at the dog, waving the stick and shouting all the way. Let's pause the episode there. I don't know how things turn out. Maybe this will actually work better than the strategy Benny judged best. But what is he thinking? Unlike Adam, Benny's eyes are on an appropriate target. He just has unjustified beliefs about what's the best way to carry it out. That's *his* mistake number 1. And then, he disregards those beliefs. Not because he suddenly saw their weaknesses. But just out of nervousness or bad planning or the ordinary practical stupidity. That's his mistake number 2. What kind of mistake was that? Was it that he was after all *justified* in the taking the slow approach, unjustified though he was in thinking so? I hesitate to say so. This case seems rather to have much the same feel as Adam's case, only now with unjustified beliefs about means substituting for unjustified ends. Yet it's not that those beliefs—beliefs Benny should have known better than to have—justified him, or gave him reasons, to pursue the means. It seems to be a different kind of pressure they impose on him, and that he's now wrongly flouting. That is, he's doing something wrong in flouting it. I don't know whether, by now flouting that pressure, he's *overall* doing better or worse. Still it does seem there's something he's flouting.

Or consider Clarissa, who thinks, against all evidence to the contrary, that her dog wouldn't hurt a fly. As it happens, there are questions about her dog's

history. He just wandered in one day a few years ago. Some neighbors wonder if he might be the dog of that old man who went to prison. They look a bit alike. But Clarissa doesn't believe it. Then she hears a story about how that old man's dog was extremely gentle. Just like she (unjustifiably) thinks her dog is. But this doesn't affect her confidence about her dog being the old man's dog one bit. I think Clarissa is being a bit unreasonable here. She shouldn't be thinking her dog is gentle, in the first place. And given her dog's evident ungentleness, it probably isn't very likely to be the old man's dog. However, Clarissa does *think* he's gentle. So what does she think she's doing, ignoring this new bit of evidence. I don't want to say that Clarissa now has more *reason* or *justification* to think her dog was the old man's. Since her evidence speaks against her dog's being gentle. But there does seem to be some kind of pressure on Clarissa to think so, which she's doing something wrong by ignoring.

Or consider Diana, who thinks (unjustifiably) she's no good judge of her neighbors' characters. In fact she has a keen eye and her suspicions are usually right on target. Her doubts aren't backed by evidence but by a crippling lack of self-worth. Still, they are doubts she does have. Then she meets a new neighbor and decides he's not to be trusted, and she acts on that belief. Ignoring her own, still held, beliefs in her own unreliability. What does she think she's doing? Alright, she shouldn't be thinking she's unreliable in the first place, but given that she does, shouldn't she be more circumspect or critical about her first reactions? I don't want to say she has *less reason* to think the neighbor is untrustworthy than she'd have if she had fewer self-doubts. For her self-doubts aren't backed by good evidence. But it still seems she's doing something wrong here. Something additionally wrong, besides having those self-doubts in the first place. There's something normatively awry about her lack of reflective coherence here.

Or consider Elena, who thinks (unjustifiably) she ought to read her daughter's email, yet she doesn't do it. Sometime later the daughter elopes, though between us, this doesn't vindicate Elena's earlier thinking. Her reasons for thinking she ought to read the email really were no good, and she wouldn't have seen anything about the elopement there anyway. Yet now Elena is kicking herself, "I knew I ought to have pried more, why didn't I do it?" I sympathize with Elena's self-criticism. She believed she ought to  $\phi$ , yet she didn't. What's up with that? Of course, her normative belief wasn't justified. So I feel some conflict here. I don't want to say that whenever you think you ought to  $\phi$ , no matter how unjustifiably, you should. At the same time, I think that in acting against her own normative beliefs, Elena also exhibits some kind of normative failing. There's some pressure here that she's flouting. But it's not the same kind of pressure as she'd be flouting if she really had reasons to read her daughter's email.

Or consider Frank, who has different issues with his children. See, Frank's friends have persuaded him that he ought to be maximally non-interfering with his children. Don't tell them when to go to bed, or anything like that. Let's suppose, for the sake of discussion, this is wrong. This is really not the way he

ought to be interacting with his children. But let's suppose also that Frank's friends have given him good reasons to think it is. So his belief about what he ought to be doing is justified, it's just false. Is this coherent? I think it is; though I know some philosophers would deny it. Maybe they'll persuade us. But for the moment I'll go with the face-value judgment that beliefs about what you ought to do can be justified yet false. That's really not the way Frank should be raising his children. It's not just that it will have a bad outcome. It's really not what he has most reason to do. But he does have good reasons to think it is. Perhaps the good reasons to think it's what he has reason to do affects the reasons he has. Perhaps it gives him more reason to raise his children that way than he would have otherwise had. I don't mean to resist any such thought. It just still seems that, even after that's taken account of, he might still on balance have not very much reason to behave that way, but substantially more reason to think he does.

Unlike the previous examples, I'm not going to add further twists about Frank's incoherent follow-through. I introduce him now just as a counter-point to Elena, that we will return to later. Notice the difference between these cases. Elena has an unjustified higher-order belief, which she flouts. Frank on the other hand is justified in his higher-order belief, but that belief nonetheless is false. You may find it easier to admit the possibility of cases like Frank if you think in degree-like terms. Some philosophers think if you're justified in thinking you're justified, you must then be justified. As I said, I am prepared to think that justification to think you're justified should have some effect on your first-order justification. But I see no reason to be sure that whenever someone has such-and-such much reason to think they're such-and-such justified, they will in fact have that much first-order justification.

We will come back to the case of Frank later.

## Section 2

I dub the kinds of "oughts" we're targeting here "hypothetical oughts." The idea behind this name is that the means/end principles we began with seem to have some kind of hypothetical force—hypothetical on what ends you actually have and what beliefs about means, not just what ends and beliefs you *ought* to have. On the other hand, other normative sources, both practical and epistemic, purport to have more categorical force. In these latter cases, you seem to really have *reasons* to act and think in certain ways, in a way that bears on you normatively differently than in the cases described above. The normative pressure in these cases seems different, and seems not so dependent, or at least, dependent in different ways, on what attitudes the agent happens then to have.

The name is also motivated by the fact that, in describing cases like the ones in Section 1, we're tempted to say things like: It's not that Adam *outright* ought to be buying poison. It seems better to say he only ought to be buying the

poison, *given* or conditional on his having the end he does. And so on *mutatis mutandis* for the other examples.

However, this name “hypothetical ought” should be regarded as a stipulated not a descriptive name. It may turn out to be an imperfect description of the phenomena.

For one thing, it’s not straightforward how tightly the phenomena we’re investigating is tied to “oughts.” There may be much that philosophers find interesting about “ought”, even more specifically about deontic “oughts”, that doesn’t necessarily belong to our intuitive target. Nor, as we’ll see, is our intuitive target restricted to “oughts.”

There are a host of slogans and ideas philosophers already associate with the kinds of cases we’re considering. They’re not entirely without merit, but on balance I’m often inclined to resist them, at least in part.

One idea is that we’re dealing here with a “wide scope” ought. It’s not that Adam ought to be buying poison. Nor even that his end implies he ought to be buying poison.—What would it be for *his end* to imply this? It’s not obviously a proposition, and even if it were, it’s not obviously a proposition that on its own has any normative consequences. It might instead be a proposition like “The dean is dead” or “I kill the dean.” The latter would at best entail that Adam *does* buy poison, not that he ought to. Should it be *Adam’s having that end* that implies he ought to be buying poison? Well, then, since he does have the end, presumably it would then be true that he ought to buy the poison. But we said we wanted to resist that consequence.

The “wide scoper”’s idea is to reject all of this. It’s not that Adam ought to buy poison, nor that anything (or at least, that anything non-normative) implies he ought to. What’s instead true is that Adam ought to be such that *if he has such-and-such an end, he buys poison*. He ought to be such that that conditional is true of him. Which is compatible with his in fact oughting not to have the end, nor to buy the poison.

The kernel of this idea has a long history. Schroeder traces it back to Hill, Darwall, . In the past thirteen years it’s been prominently defended by and associated with John Broome. I’m *broadly* sympathetic with much in this tradition. But many distinctive elements of Broome’s view will turn out, from where I end up, to look sometimes idiosyncratic and sometimes unpalatable. So while in a way the position I advocate is broadly Broomean, you may or may not find that association helpful.

Another idea in the vicinity, one that plays a large role in Broome’s own presentations, is that we’re dealing here with failures and requirements of *Rationality*, rather than with Reasons and Justification. These claims don’t obviously seem right to me. I’m willing to go along with them as stipulations of a technical notion of “Rationality,” but I doubt that’s what they’re supposed to be. I don’t know *what all* should and shouldn’t be counted as Rationality, in the

pre-theoretic understandings of that term we bring to the table. But it's not obvious to me that the notion of Rationality exactly lines up with the phenomena we're considering. If it does line up, I'd prefer we discover that to be so as a substantive result, rather than build it into our specification of what phenomena it is we're investigating.

Part of the idea of calling the phenomena we're talking about "requirements of Rationality" and separating them from "requirements of Reason" is that the kind of normative pressure imposed on Adam by his merely intending to kill his dean, in abstraction from whether he's justified in doing so, need not amount to his having justification to kill his dean. With this I am in complete agreement, and our initial series of examples mostly tried to evoke that thought from you, as well. But another element in existing discussions of this is a particular conception of Reasons. That is the conception that normativity is somehow external and may be more objective than many philosophers think, that it may radically outstrip what agents then have access to. This is the conception at use when someone says the person who's just been served poison at the cocktail party has a Reason not to drink from their glass, though there may be no evidence available to them which says so. This conception of Reason may or may not be correct. Its analogue in the epistemic domain would be even more controversial than it is in the practical domain, though there too, it may or may not be correct. Still, this issue seems to me to be orthogonal to the phenomena I've targeted. You should want a distinction between what I'm calling hypothetical and categorical normative pressures even if you're a constructivist or a Humean. It's only a certain radical kind of subjectivism that would see no need for this distinction.

Because these debates tend to get mixed together in existing discussions, and because I'm not sure what our pre-theoretic notion of Rationality amounts to, anyway, I will abstain from saying that we're talking here about Rationality as opposed to Reasons or Justification. With the initial idea that Adam isn't justified in buying poison, though, I am as I said in complete agreement.

A third idea in the vicinity is that my examples demonstrate "subjective oughts" and "subjective rationality" rather than their objective cousins. When people use those labels, they usually do have in mind the same phenomena we're here investigating. But as we'll see, I think these labels encourage substantive assumptions about the phenomena that (1) I don't want to build into our specification of what it is we're talking about, and (2) I think are probably false. So I discourage the use of these labels too.

(A thought guiding me here is that the "oughts" on the hypothetical side might not prove to be completely transparent. They are presumably *affected by* what you think you ought to do, whether justifiedly or not. That was one of the intuitive motivations for making the contrast. But their being so affected doesn't mean they must necessarily match what you think you ought to do. I think of the phenomena we're investigating as something like a salad, and your beliefs about what you ought to do as being the cucumbers. Some salads are made wholly

from cucumbers, but other times they're just one ingredient among others; and you can also have salads where they are absent.)

The oughts in our examples are often talked about as “oughts of advice.” It's not that we really think that Adam is normatively best-served by buying poison. He's best-served by giving up both his assassination project and the poison-buying that would advance it. But given the ends Adam does have, we might advise him to buy the poison. At least, we might do so if we're in a certain frame of mind. In other frames of mind, we'd advise him not to murder. Or if his ears are shut to such advice, we might advise him to ask all his neighbors if they have any spare poison—since that may get him apprehended before anything worse happens.

What is really going on, when we would advise Adam to buy poison today? Are we thinking, it may be bad for him to be an assassin, but it would be worse for him to be practically incoherent? Of course we don't think that.

I really don't have a good understanding of the range of things we try to do when we give advice, or what the rules are. We're not always encouraging what we believe to be the best outcome, though sometimes we are. We're not always encouraging what we think the agent is categorically most justified in doing or thinking, though sometimes we are. We're not always putting ourselves in the agent's shoes and giving voice to what we think is the most sympathetic and competent downstream choice, though sometimes we do that too. Sometimes we seem to be doing mixes of all these things at once. I acutely feel the lack of a good taxonomy of the kinds of aims we have and rules we play by, across the range of activities we informally call “giving advice.” Perhaps once this is all critically better-organized, we'd be able to identify one of the modes of advice giving, in terms independent of the phenomena we're now investigating. And it might be that our phenomena turns out to closely track, or even be more intimately tied to, what we do when we give that kind of advice. At present, though, I don't think we have any such *clear and independent* grasp of the kind of advice-giving that verifies such a connection.

Once I tried to develop an extended example, involving two students hitchhiking in a foreign country, about to buy a present to influence an authority, but disagreeing about some of the empirical and normative facts about their situation. Holding all that fixed, I tried to imagine the one student advising the other in various ways, in different advice-giving spirits. How many different recommendations could I intelligibly come up with for him to offer? I got up into the teens, or at least persuaded myself I had, before giving the exercise up. I wasn't seeing any systematic principles organizing the variations. And many of them required attention to subtle details and differences I wasn't optimistic about getting others to exercise, for little discernible payoff. But that's the result of my most considered guesswork about how many different ways there are to give advice. At least some ten many.

Now there are some unification strategies that might be deployed here, and



reduce the number of options. Think of Lewis on “can”: Holding certain facts fixed, you can raise your arm right now, though you do not. Holding other facts fixed, you cannot. Some formal treatments of “ought” take a similar approach. Such views do help us make progress understanding the phenomena. But I have doubts about some of the constraints they impose, and distinctions they seem ill-suited to explain. Without getting into details, I think we do best to work through the informal organization I attempt below before bringing in any existing or new formalism.

### Section 3

It will help now to mention two famous kinds of puzzle cases that are often associated, more or less loosely, with the phenomena we’re considering. In my view, while both are interesting, neither obviously bears on the very phenomena I’m targeting.

The first is examples like Parfit’s miner. In these examples, you have two options, one of which will lead to a terrific outcome and the other to a terrible outcome. You don’t know which leads to which. You also have a third option, which you know leads to a merely good outcome, one which is clearly inferior to the terrific outcome. For example, some miners are all trapped in tunnel A or in tunnel B, you don’t know which. Water is coming to flood the tunnels. You have the opportunity to divert all the water to tunnel A, or to divert all the water to tunnel B, or to let the split water between the tunnels. In the last case, you know some of the miners will die but the rest will survive. In the former cases, if you choose luckily, all the miners will survive, but if unluckily, they’ll all die.

In these examples, many want to say you ought to, or are most justified in, choosing the third option, despite knowing it gives a sub-optimal result. At the same time, people are tempted to say things like “Well, if or given that the miners are in tunnel A, then you ought to send the water to other tunnel. And likewise if they are in tunnel B.” But since you don’t know which tunnel they’re in, we refrain from concluding that there’s some (unknown) specific tunnel down which you ought to send all the water. This is interesting, because the miners are all either in tunnel A or in tunnel B. So it looks like we want to say, “if A then you ought to X (send all the water down a specific tunnel); and if B then you ought also to X (albeit in this case, you X in a different way); and A or B is true; but it’s not the case that you ought to X.” MacFarlane and Kolodny have discussed the puzzle in this form, arguing that it provides counter-examples to *modus ponens* different from McGee’s famous counter-example.

Now the phenomena discussed here *may* be continuous with the issues we’re discussing, if you think that what you ought to do may be affected by things unknown to you. All of us sometimes informally talk as though it can, but when we get down to rigorous theorizing, many shy away from such possibilities.

Someone who shies away from saying Williams' about-to-be-poisoned drinker at a party ought not to drink may also shy from saying we ought to send all the water down tunnel B, even if the miners all in fact do happen to be in tunnel A. They'd try to find some other theoretically careful way of describing this case. Given this, they may reject that:

Given that the miners are all in tunnel A, you ought to send all the water down tunnel B.

is true on the same understanding that:

Given that he intends to kill his dean, Adam ought to buy poison today.

is—whatever the latter ultimately amounts to. I don't insist that these absolutely do come apart, only that on many ways of carving up the normative, they will. So we can't start off by assuming that these two sets of issues are continuous.

One difficulty with the linguistic treatments of "ought" I know about is that they don't draw distinctions here. And perhaps there is no linguistic marking for the distinction I've here in mind; I don't know. But most normative theorists will see fundamental differences between what's going on in the miner examples and what's going on with Adam and the other examples I started with. This is one reason it's good not too much specific focus on all the specific wriggles in our informal usage of "ought." More reasons will be listed later.

Another interesting thing demonstrated in these examples is that the normatively correct option may be one you know to have a less-preferred outcome than other options. I do think that is a genuine possibility, and may be present in some epistemic cases as well. For example, it may be normatively correct for you to simultaneously believe each of a set of claims, while knowing them to be jointly inconsistent. (You just aren't in a position to sort out which of them are false.)

The second example is Jackson's Professor Procrastinator. He's invited to review a paper. He knows that the ideal outcome would be for him to accept the invitation and complete the review in the near future. However, he also knows that he's unlikely to complete the review soon, even if he does accept the invitation and agrees that he ought so to complete it. He knows he's a procrastinator, and unlikely to reform. So given this, we're invited to judge that the Professor ought instead to turn the review down.

That is, of the three actions:

Accept the invitation and complete the review quickly.

Decline the invitation.

Accept the invitation and complete the review more tardily, if at all.

the Professor ought most to perform the first, but knows he won't, even if he accepts the invitation. The second option is inferior to the first but superior to the third. So between accepting and declining, the professor ought to decline. Or rather, he ought to decline, given that he's a procrastinator. In some more categorical sense, he ought not to be a procrastinator in the first place, and also then accept the invitation rather than decline.

This example also raises interesting issues about the logic of "ought." On some ways of working all this out, we might say that the Professor ought to X and Y without its being true that he ought to X. He ought to accept and complete the review quickly, but ought not to accept simpliciter, since if he does, he won't complete the review quickly. On other ways of working it out, we would still insist he ought to accept simpliciter, rather than decline, but say instead that he ought to decline given that he's a procrastinator. Rather than that he ought to decline simpliciter.

There isn't a consensus on how we should analyze and talk about such cases. But they do look suggestively similar to some of the things we find it natural to say about Adam and the other characters we began with. We shy from saying that Adam ought to buy poison today, simpliciter. But rather want to say things like, given that Adam intends to kill his dean, and poison is the best or only way to do it, he ought to buy poison today.

However, despite some similarities in how we express ourselves, there is a fundamental and important difference between our earlier examples and the present one. In brief, the kind of dependency or hypotheticalness in the Adam cases is upon the *oughtiness* of the ends he happens to have (perhaps not oughtily). In the Jackson case, the dependency is just upon the *presence* of the Professor's ends and existing character. I can explain this better with the following sets of examples. They come in two groups of three.

In the first trio of examples, it's common that you believe unreasonably that it will rain today. Your evidence says the day will be nice, but you perversely believe against the evidence. Now in the first fleshing-out of this common background, it's also true that God has decreed that anyone believes unreasonably that it will rain and carries a heavy bowling ball will get a reward. Then it does seem that given your unjustified belief about the rain, you ought to carry a heavy bowling ball. Of course without that unjustified belief, you shouldn't, because the ball is heavy and no good will come from carrying it. In the second example, God has decreed that anyone who believes it will rain, simpliciter—whether justifiably or not—and carries a heavy bowling ball will get the reward. Here, it does seem that given your belief about the rain, which happens to be unjustified, you ought carry the bowling ball. The ought here is insensitive to whether your weather belief is justified. It comes just from the fact that you have it. In the third example, God has decreed nothing, but you still have your unjustified belief about the rain. In this case, we may want to say, given your belief about the weather, you ought to carry an umbrella. But it's not true that you ought to carry the umbrella simpliciter, because you really don't have good

reasons to expect rain.

That is, in all three examples we'd say something like this:

Given your unjustified belief about the rain (and its unjustifiedness), you ought to X.

Given your belief about the rain (which happens, irrelevantly, to be unjustified), you ought to X.

Given your belief about the rain, you ought to X.

but the normative structure of the cases is interestingly different. In the first two cases, I think it's categorically true that you ought to carry the bowling ball. In the last case, I wouldn't say you categorically ought to carry the umbrella. All that's categorically true in the last case is that you ought to believe otherwise than you do. (This is also true in the first two cases. Though you may end up better off with this unjustified belief, and getting God's reward, than you would merely by believing more carefully.)

That's the contrast I tried to express briefly by saying in some cases, the ought depends on the mere presence or your actual attitudes. In the last case, the dependence seems more tied to the (perhaps absent) "oughtiness" of your upstream attitudes. In a way, it's because of the hypothetical oughtiness of your rain belief that you ought, in whatever sense you ought, to carry the umbrella. But in the second case, merely by virtue of *having* the beliefs about rain, you already ought to carry the bowling ball. In the first case, perhaps the dependence is hypothetical not on the presence or the "oughtiness" of your beliefs, but rather on their "non-oughtiness."

In the bowling-ball cases, we're talking about the bearing your irrational beliefs have on you qua facts (perhaps including the fact that they are irrational). In the umbrella case, we're talking about the bearing they have on you qua upstream input to your normative faculties.

The phenomena I want us to investigate are illustrated by the umbrella example. Jackson's Professor, I think, is more like the bowling ball examples. In his case, some of the mere facts about what the Professor's ends and intellectual dispositions are make a difference to what he ought to do, a different difference than would be made if he had normatively better ends and dispositions.

Now there are interesting questions about when you are entitled to treat some of your own attitudes and dispositions as mere facts in the world, that relative to a given deliberation are just there to be worked around, rather than things you have to right now actively try to control. When is the Professor entitled to say, "Look I'm a procrastinator, so I ought to decline this invitation." And when is he required instead to say, "Look I've got to refrain from procrastinating here." Can saying the first co-exist with also saying the second? Or do they displace each other?

Those are interesting and important questions. But I think they're somewhat different than our present topic. Or at any rate, there's a lot of interesting things for us to settle even if we bracket the issues raised by Jackson's Professor. Our focus is on what is the downstream normative bearing of your attitudes qua deliberative inputs, when it's bracketed whether those attitudes are reasonable.

(There's also a famous contrast, but I don't know offhand where it originated: if you think everyone's out to get you, you ought to check yourself into an institution. That's Jackson's "ought." Or is it that: you ought then especially to *avoid* people in white coats. That's Adam's "ought.")

You may have the thought that Jackson's case has to do with *the fact that* you have some attitudes, and what difference that makes. Whereas our target has to do with *the contents of* your attitudes, and what difference they make, or purport to make. However, this is not right, for several reasons.

First, some of the phenomena we're considering have essentially to do not just with the content of your attitudes, but with the fact that you have those attitudes. It's essential to the defect Adam displays when he fails to buy poison that he does in fact intend to kill his dean.

Second, it may be better to talk about *the downstream inferential impact* of Adam's actual attitudes. I would resist identifying this with what follows from *the contents of* his attitudes. As I argued at the end of my Anti-logic paper, which we read at the January conference, it's doubtful that downstream inferential impact is confined solely to relations at the level of your attitudes' contents. To remind you of one of the examples I discussed there, your withholding belief in P will tend to (or ought to) have a downstream inferential impact. But it's not clear what is the content from which that impact is supposed to derive. Talking about downstream inferential impacts may be a way of bracketing that debate, about whether this is just a matter of contents.

However, I'm reluctant to think that the phenomena we're discussing here arise only at the level of psychologically real *inference* and deliberation. So it'd be best to work with a notion of "downstream impact" which need not be psychologically manifest, and where we also are bracketing the question of whether the upstream impacts are really normatively appropriate. It's like talking about what your evidence supports, regardless of what you consciously think. Only now we're talking about what your actual attitudes support, regardless of what you occurrently judge. And regardless also of whether those original attitudes were reasonable.

I don't think we have a precise pre-existing theoretical vocabulary for discussing these kinds of "downstream impacts." At any rate, not one that we all already understand in the same way. That's what I'm trying here to carve out.

If I understand Christian properly, he would call this the "is-a-reason-for" relation between attitudes. Your actual belief that it will rain is-a-reason-for bringing an umbrella, though you're not justified in bringing an umbrella since

your belief in rain is one you ultimately have no good reasons for. I think this vocabulary is too easily misinterpreted, so prefer not to use it. But I think that what he's referring to thereby is the same kind of normative relationship that I'm talking about with "downstream impacts" and "hypothetical oughts."

I promised a second trio of examples. These illustrate the same contrast as the bowling-ball-and-umbrella trio, but may help by doing so from a different angle.

What's common to the examples is: You are guessing how many jelly beans are in a jar, and your first guess is that there are 1000. In fact, you have evidence that you tend to err on the high side in these guesses, and so your evidence best supports a final estimate of 900. However, as it happens, you believe, contrary to evidence, that you tend to err on the low side in your guesses. If you followed through on that belief, you ought then to arrive at a final estimate of 1200. But in fact you ignore both your undermining evidence and your undermining belief, and stick with your initial guess. You believe there are 1000 jellybeans. This belief is epistemically defective, though settling whether and in what senses you ought to revise it upwards or downwards is tricky. I will propose machinery to describe this case more carefully later.

Also common to the examples is that you confront a series of doors that are labeled "800 jellybeans", "900 jellybeans", "1000 jellybeans", and so on. God has put a prize behind one of the doors, not necessarily the one that reports the actual number of jellybeans. And as a matter of fact, there are 600 jellybeans in the jar. It only looks like there are 900.

Got all that? It's a complex set-up, but its complexity enables us to illustrate a variety of different possibilities. Now we will flesh the case out in some more specific ways, corresponding loosely to the three ways of fleshing out the bowling-ball-and-umbrella examples.

In the first fleshing-out, God won't put the prize behind the door with the correct estimate. Rather, he will tend to put it behind a randomly chosen door. But, he will note what a contestant's actual estimate is, and whether that estimate was formed irresponsibly. If it was, he will in that case put the prize midway between the actually correct estimate and the one you irresponsibly hold. (This is a fickle God.) Given the facts as I described them, then, in this case the prize will be behind door 800. If you had a more reasonable estimate, or even (unreasonably) had the correct estimate, the prize might then be anywhere.

Some may want to say: given all these facts, including your irresponsible belief that there are 1000 jellybeans, you ought to look behind door 800. Though many will also shy from saying this, especially those who shy from saying Williams' about-to-be-poisoned drinker at the party ought not to drink. Door 800 is, at least, the door that can be expected, objectively, to give the best outcome.

In the second case, God decrees that the prize will be midway between wherever you *do* believe is the right number of jellybeans and their actual number—regardless of whether your belief is responsible or justified. Here too, given your

actual belief, looking behind door number 800 can be, *objectively*, expected to give the best outcome.

(What *you* ought to expect, though, is trickier. Ought you to expect door 1000, since that's both what you know yourself to believe is the number of jellybeans and what you believe is the correct estimate? Or ought you to listen to your actual belief that you guess too low, and so believe that there are 1200 jellybeans—so perhaps the prize is behind door 1100? Or ought you to believe there are 900, since that's what your total evidence in fact best supports. There's a great amount of variability in what answers you can elicit here.)

In the final case, the prize is just behind the door that correctly reports the number of jellybeans. (The case you were waiting for!) Here looking behind door 600 will actually give the best outcome. Perhaps you ought to expect the prize to be behind door 900, since that's what your evidence really testifies is the right number of jellybeans. But since you actually believe 1000, we are also tempted to say:

Given that you believe there are 1000 jellybeans, you ought to think the prize is behind door 1000.

This seems to be the same kind of hypothetical normative judgment that is also elicited by Adam's case, and the other examples from Section 1. The kinds of dependence on your actual belief illustrated by the earlier members of this series are different.

The examples in this series have more moving parts than the bowling-ball-and-umbrella example, and so they are more difficult to think about. They do illustrate though, that our judgments of "Given he thinks this, he ought to do that." are not always confined to tracing out what actions would be consistent or inferentially downstream from what the person thinks. But the judgments that track our target phenomena do seem to loosely respect a constraint of that sort. (I say only "loosely" because as we've already remarked, our target phenomena need not coincide with consistency or any relations solely between contents. But if we loosen up to admit a notion of what's downstream from some starting premises that may include conclusions you know to be inconsistent, but still best supported by those premises, that's the notion we want here.)

This series of examples also illustrates that the kinds of hypothetical normative pressures we're thinking about can come from unjustified undermining beliefs, as well as unjustified supporting or opposing beliefs. Those are kinds of hypothetical normative pressures I've relied on in earlier discussions of these issues. (In papers about perceptual skepticism and Moore's Proof of an external world.)

## Section 4

Let's say a few words about "ought."

First off, I observe that many earlier discussions of our phenomena, and related issues, are treating “ought” as a kind of “must”—in our cases, as more specifically a deontic must. However, as some linguists have argued (at least von Fintel), “ought” really is weaker in natural language than “must” or “is required to.” It makes perfect sense to say, in English:

You ought to go to the party, but you don’t have to.

or, just as easily:

You needn’t go to the party, but you ought to.

However, neither of these makes sense with “must” substituted for “ought.” Of course, there is some latitude in the diagnosis of such cases, because different sources of normative pressure might at once be contextually salient. But my understanding is that there are robust patterns here, and that “ought” really does behave in a way interestingly weaker from “in all required situations, you do...”

However, all philosophical discussions of “ought” that I now recall have been insensitive to this. I will go along understanding “ought” to mean what philosophers treated it as meaning, not as what it actually does, in English mean—whatever that is.

Second off, I observe that the phenomena we’re discussing aren’t restricted to cases where the normative pressure you’re under is a *requirement*. In some of his early papers, Broome acknowledged that we might talk both about normative requirements and normative recommendations, where these were the wide-scope counterparts of all-things-considered and merely pro-tanto reasons. In his later work, he changed his mind about using the term “normative” to describe the wide-scope notions, since that pre-judged that we had some reason to conform to them, and that he came to regard as a substantive further question, one whose answer he came to suspect might even be no. So in the terms he later preferred, these would be called *rational* requirements and recommendations. Whether they had any normativity to them turned on whether we had reasons to conform to them. I mentioned before my reservations about using Rationality-versus-Reasons as the way to mark the operative contrast here. I mentioned also that I think of wide-scoping as just one particular way of analyzing the phenomena. Our grasp of the phenomenon itself shouldn’t be tied to that specific analysis of it. But I want to call attention, in this sketch of Broome’s views, to the existence of *recommendation-force* pressures of the “merely rational” or “merely hypothetical” variety, as well as *requirement-force* pressures.

I haven’t looked yet at Broome’s most recent work, but for a while now in his work the recommendation-force pressures have seldomly appeared, and never seemed very important. However, to my mind they are the most obvious examples of our phenomena. Is Adam really *required*, even hypothetically, to buy



poison today? Perhaps it's merely what would best suit his plans, and isn't absolutely essential to them. And so on for our other examples. Broome restricts his attention only to the cases where it's most plausible that there is some such rational-or-hypothetical requirement. I'm not sure that there are any normative requirements, though, even on the categorical side. It's at least as unclear whether they exist on the hypothetical side. For all I know, all normative pressures, of either sort, may merely be *pro tanto*.

As well as recommendation-force pressures, I think there are also, plausibly, undermining-force pressures of the merely rational-or-hypothetical sort, too. This was illustrated by our example of Diana, who has an unjustified undermining belief. She thinks, against the evidence, that she's no good judge of her neighbors' characters. Then she manifests some kind of defect in ignoring this belief when she assesses a new neighbor's trustworthiness.

Hence, in my thinking the phenomena we're discussing display a whole range of "forces": sometimes supporting a belief or action, sometimes opposing it, sometimes undermining, perhaps sometimes calling for withholding. Perhaps, though it's controversial, sometimes requiring beliefs or actions in a way that's more than merely *pro tanto*, and cannot be defeated. Broome mostly only talks about our phenomenon in this last form.

There is a reason for this, having to do with one of his central motivations for talking about merely-rational-or-hypothetical pressures in the first place. Broome does offer several arguments for the package of views he proposes, adding and emphasizing different considerations in different papers. But his most prominent and consistent argument goes like this.

Consider someone who is practically akratic. Someone who is epistemically akratic may serve as well. These agents believe, without any antecedently good reasons, that they ought to  $\phi$ . If they then fail to  $\phi$ , we think something has gone wrong. But what?

The simplest, subjectivist view here would say that whenever you do in fact believe you ought to  $\phi$ , it then becomes true that you ought to  $\phi$ —even if it wasn't true before you thought so. Most of us shy from this position. As Broome often describes it, it seems to legitimize an objectionable kind of normative "bootstrapping." Think whatever you like about what you ought to do, because those beliefs will then be true.

Broome assumes all parties to our inquiry will want to avoid that result. Then he considers two competing ways to go. One view would continue to say that your actual normative beliefs, even if based on no antecedently good reasons, still do give you *some* reason-like support for doing as they envisage. It's just that this reason-like support may also be opposed and ultimately outweighed by other reasons you have to do otherwise. In my vocabulary, these views say that actual normative beliefs, even if antecedently unjustified, do generate some normative pressure of the categorical, genuine-reasons-and-justification sort. It's

just pressure whose force is merely pro-tanto and recommendation-like, and can be defeated.

Broome's alternative is to say that your normative beliefs continue to generate pressures whose force is requirement-like, and infeasible. It's just that the kind of pressure they generate is merely wide scope, or more neutrally, hypothetical, rather than of the categorical, genuine-reasons-and-justification-to  $\phi$  sort. (In Broome's picture, wide-scope requirements to  $\phi$  given  $\eta$  may turn out to involve categorical reasons to do something, but it won't be to  $\phi$ . And it's not obvious that they ultimately do turn out to involve categorical reasons to do anything.)

Broome thinks his alternative has a definite advantage here, in that he can continue to say, while his opponents cannot, that the kinds of pressure involved are not merely pro tanto and defeasible. As he puts it, if you think you ought to  $\phi$ , and you don't, then regardless of whether your normative belief was antecedently justified, *there is definitely something you're doing wrong*. It couldn't be that you merely have some pro tanto considerations weighing for  $\phi$ -ing, but might on balance be doing best, since other pro tanto considerations of the same sort might weigh more heavily for doing something else.

Broome emphasizes and insists upon this argument repeatedly. It seems to have played a central role in his thinking about the phenomena. As I said, I haven't yet read his most recent book manuscript, so I don't know whether it still has that central role.

I said already that I'm not sure any normative pressures, of either the hypothetical or categorical sort, have the force of infeasible requirements. And in particular, I doubt that normative beliefs do. I think that thinking you ought to  $\phi$  merely does offer some consideration for  $\phi$ -ing, which may not always on balance win out. I mentioned an example relevant to this in our opening section—the final example, of Frank. We will come back to that case momentarily.

So I'm not myself motivated to separate hypothetical from categorical normative pressures for (all) the same reasons Broome is.

I am motivated instead by its seeming to be necessary, to make the distinctions we intuitively want to draw between our opening examples, and cases where subjects really are justified, and really do have reasons, to form the same beliefs and perform the same actions.

Also I am motivated by the conviction that having a clear grasp of these distinctions is important for steering more surely in other, more distantly related, philosophical inquiries. To take one example, there has been much discussion this past decade of "the problem of disagreement." Suppose you and another agent know yourselves to have access to all the same evidence about some question, and you would have antecedently judged the other as reliable an assessor of the evidence as you are, but now you reach your independent verdicts and find they don't agree. What weight should you give the fact of the other's disagreement in your final estimate of the answer? No weight—you've already factored

in all the evidence they were relying on? Some weight—her disagreement makes it more likely you’ve screwed up, still the evidence does seem on balance to point somewhat in the direction you supposed. Or “equal” weight—you should still count each of you as in some sense equally likely to be right? Or do these not exhaust the interesting possibilities. Tom Kelly has argued that all of these options go wrong in treating the issue as though it were merely a “formal” one, which could be decided by the mere fact that there are these disagreeing parties, without attending to what the original evidence really does say and what estimate each party originally formed. Kelly asks us to consider the case of Right and Wrong, who are as we’ve already described, except that Right’s independent verdict really did conform to their shared evidence and Wrong’s didn’t. Or the case of Wrong and Wronger, who both interpret the evidence badly, but one more egregiously so. Do we really want to say that in both cases, the two parties’ epistemic situation is completely symmetrical?

One controversial issue is whether Right should revise his estimate towards Wrong’s at all. A strong form of the claim Yes would say that Right should revise his estimate to a degree that is in some natural sense “halfway” between his original verdict and Wrong’s. (It is of course difficult to say what “equal weight” and “halfway” amount to, in a rigorous way. But I think that is a distraction from the most interesting issues here.) Even if one did think that Right should move some of the way or “half” of the way towards Wrong, though, one might have additional qualms about saying that if Wrong made the corresponding move towards Right, he would then have the view he’s (categorically) justified in having. Why should halfway between his first crazy view, and what his evidence really supports, now be what he’s justified in believing? If he never heard about Right’s verdict, would he have been justified in sticking to his first, crazy view—despite its being contrary to his evidence?

And what about Wrong and Wronger? They start off with views that don’t fit the evidence. If they revise *towards* each other, and end up with views that fit the evidence *worse* than the better of the two did formerly, are they then epistemically better off?

Kelly urges that we instead recognize that it’s not just the abstract shape of the disagreements that matters. It also matters what the evidence really supports, in the first place, and whose opinion more closely matches that.

David Christensen, who defends the “equal weight” style of views, acknowledges this point, and says that we should understand arguments for the answer of “some weight” or “equal weight” to be arguments in the first place about what it’s right for agents like Right to think. What it’s overall most reasonable for agents like Wrong or Wronger to think will be more complicated, and will be sensitive to the real force of the original evidence, which they’ve mis-assessed, as well as to the facts about their differing verdicts.

That is a good strategy for bracketing issues about the normative pressures of unreasonable deliberative inputs. But I find it noteworthy that it’s taken about

a decade for the discussion to get clear about this. (I emphasized the point in discussions and seminars, and some other authors have said things along these same lines. But my sense is that it's only in the past year or so that one can start to take appreciation of this point for granted.)

A different thing we might say is that, at the level of hypothetical normative pressure, the agents *are* all in symmetrical positions. It's only when we think about what categorical difference that should make that things get tricky. (What I've said here is not obviously true, because it assumes that the differences in how the agent's actual attitudes fit their evidence didn't leak into their hypothetical normative positions. That is, it treats hypothetical normativity as though it were an entirely a matter of subjective coherence, independent of where the evidence objectively does point. This may be the right way to go, but it's not obvious that it is. We'll mention this issue again later, without trying to settle it. I don't claim that the idea that the disagreeing subjects are in symmetrical positions at the level of hypothetical epistemology is definitely correct, just that it's a natural idea to have. I think that in practice, most earlier discussions of the problem of disagreement were proceeding as though this were so, and were pursuing the question of what *hypothetical* impact the news about a peer's disagreement should have. They just hadn't explicitly articulated that this is what they were doing.)

Of course, in any such debate, perhaps there will be a strategy like Christensen describes for bracketing off questions of hypothetical normativity—questions of what kind of normative impact comes from unreasonable deliberative inputs. But until we've framed things that way, we need to be aware of the way that our intuitions and debates can be shaped by assumptions about hypothetical normativity. Assumptions that we may, unknowingly, not all agree about.

Best is to have a good understanding of what hypothetical normative pressures look like, when they're showing up in other issues we're exploring, and what the range of views and open questions about hypothetical normativity are.

It's this, rather than certainty that in certain cases agents must be violating a strict, not merely pro tanto, requirement of rationality, that makes me think the notion of hypothetical normative pressure is philosophically important.

I've mentioned Broome's main motive for focusing on wide-scope or hypothetical "oughts," as opposed to merely pro tanto recommendations. I've explained why I don't have the same focus. So the upshot of what I've been urging is that it's not only "ought"s that matter, for the general phenomena we're targeting.

Of course, as we saw in Section 3, it's not *every* "ought" that bears on these phenomena, either—not even every deontic "ought."

A final point I want to make in this vicinity is that there are questions about the formal or grammatical structure of "ought" that, though they are important and have been inadequately attended to, press less heavily when we realize that the phenomena we're discussing ranges well beyond anything we say using "ought."


I'm not sure we have any pre-existing form of expression distinctively suited to describing the range of phenomena in our opening examples. I think we do recognize a natural class of phenomena there. And as theorists, we can introduce forms of expression specially pressed into referring to those phenomena. But I'm not sure that any pre-existing bit of natural language already does that. So the connection between the grammatical details of our actual discourse and the phenomena we're thinking about are not straightforward. I don't claim they're non-existent: the ways we do in fact habitually describe the phenomena may reflect or shape the ways we understand it. But I do think the connections are not straightforward.

With that in mind, we can, however, observe that it's usual in philosophical and logical treatments of "ought" to proceed as though "ought" operates on propositions. Some theorists argue that "ought" also takes a conditional form, or even that it fundamentally takes that form, much like probability also takes a conditional form. But even on those views, it's usually assumed that the "consequent-like" argument of "ought" is still a proposition, that is, still the sort of thing that might be true or false.

And this assumption is not obviously correct. Schroeder has a nice paper detailing some grounds for doubt about this.


One thing that bothers him is that, if "ought" operates on propositions, then it should be meaningful to say such things as that you ought that God exists, or that you ought that the sun burns. And he can't understand what those claims even say. It's not just that they're false, but he finds them uninterpretable. I share his puzzlement about such claims, but I demur from thinking this shows that "ought" doesn't operate on propositions. The claim that "ideas sleep" is similarly difficult to interpret, but it doesn't follow that "sleep" doesn't operate on nouns.

Most of Schroeder's discussion, however, surveys a number of detailed linguistic arguments for and against thinking that "ought" operates on propositions. And these are arguments from real linguists, not just from philosophers playing at it. The evidence is far from univocal, and on balance, it's at least not obvious that "ought" does operate on propositions. The salient alternative is instead that "ought" operates on a kind of linguistic element than somehow represents an action type. If so, the basic grammatical form of "You ought to swim" is not:



but instead something like:

Ought



(There may be *other* uses of ought, as in “You ought to be more widely admired,” where the “ought” does have a better claim to operate on propositions. But those uses just seem to have to do with what state of affairs is evaluatively best, unlike the genuinely deontic “oughts” we’re focusing on.)

Why does all this matter? Because if “ought” doesn’t have the grammatical form of operating on propositions, it’s not clear what claims like:

You ought: (If you intend to kill the dean, buy poison)

where “ought” purports to operate on a conditional, could mean. The conditional isn’t itself a possible kind of action. Of course there may be actions in the vicinity, and perhaps that’s what philosophers who say the “ought” takes wide scope over the conditional ultimately mean to be saying have “oughtiness.” But if that’s so, we need to hear exactly what their proposal amounts to.

Now those who treat “ought” as operating on propositions are aware that sometimes it’s grammatically awkward to say the kind of things they want to say. So they usually resort for treating “You ought that P” as meaning what would be more colloquially expressed as:

You ought to see to it that P.

Of course, “to see to it that P” is not obviously the expression of a proposition, in the same way that “to swim” is not obviously the expression of a proposition in “You ought to swim.” But these philosophers’ idea is not that they are agreeing that “ought” doesn’t operate on propositions. Rather they intend to be saying what the application of “ought” to an arbitrary proposition means.

If this were proposal were correct, then we should expect there to be just two positions where other logical elements could be inserted. If I want to stick a negation into “You ought that P,” and it has the logical structure being proposed, I should be able to stick it into either of these positions:

Not: (You ought that P)  
 You ought that: (Not P)

—or perhaps other positions embedded inside P, which we can ignore. However, when you look at “You ought to see to it that P,” there seem to be a richer range of options. It can be the case that:

Not: (You ought to see to it that P)

as for instance, when you are permitted to do anything about P. It can be of course be the case that:

You ought to see to it that: (Not P)

But there is at least one further option, namely that:

You ought: Not (To see to it that P)

I remember the following example from a survey deontic logic. It concerns my normative relationship to the disciplining of your children. The first of these claims is true: it's not true that I ought to see to it that they're disciplined. That's not my duty. Neither is true the claim that I ought to see to it that they're not disciplined. That's not my duty either. But it also seems to it that this latter is not even my right. So something more specific than the first claim can be said, and here seems to be justly said. Namely, that I ought not to see to that they're disciplined; neither ought I to see to it that they're not disciplined.

There's much to argue about here. But these observations at least raise doubts about whether "You ought that P" really can be claimed to have the logical structure more explicitly represented by "You ought to see to it that P." The latter seems to have a richer logical structure.

So, again, it remains unclear whether it really is legitimate to treat "ought" in our native languages, or our native, pre-deontic-logic thinking, as taking propositions as arguments. It may take entities of other logical types as arguments, and that may affect the range of plausible analyses of what we say about the examples we're investigating.

I agree with all of these concerns. Still, as I said, I think our ability to recognize a natural class of phenomena here isn't tied to "ought" or any other specific form of expression. So I think it should be possible to make progress in understanding them, even without settling the concerns raised here.

For more continuity with existing discussions, I will in what follows go along with the ordinary assumption that "ought" does operate on propositions, or that, if it's a binary operator in its conditional form, its "consequent-like" argument is a proposition. As I said, I have doubts whether these claims are correct, but I don't want to insist upon them.

## Section 5

Let's return to the specific proposal Broome and others have advocated, that in cases like Adam's what's true is not that:

If Adam intends to kill the dean, it's true that he ought:  
(Today to buy poison).

Rather what's true is that:

Adam ought: (If he intends to kill the dean, today buy poison).

Where that's to be understood as meaning something like that Adam ought to see to it that the conditional is true.

As I mentioned, if "ought" doesn't take propositions as arguments, it's not clear what the notion of "wide scope" here might amount to. Perhaps we can make sense of it, but it's at least unclear for now. It will depend on what the logical form of the arguments that "ought" really does take is. But we are going to bracket this line of concern.

Another worry here is what kind of conditional is involved. Presumably not a strict logical entailment: it's not Adam's duty to see to it that some of those are true rather than others. But neither does it seem adequate to say that it's merely a material conditional. At least, there are problems *equating* the phenomena we're investigating with the holding of an "ought" over a material conditional. The latter is too easily achieved—or at least, it is if "oughts" are closed under logical entailment, but that is another doubtful issue.

Broome does usually work with "oughts" taking wide scope over material conditionals, but he explicitly denies that this is constitutive of the phenomena he's discussing. Instead he calls it "the logical factor" of that phenomena. I think we understand him to be saying that what we intuitively express by:

Given he intends to kill the dean, Adam ought to buy poison today.

entails, but is not equivalent to:

Adam ought: (He intends to kill the dean \horseshoe he buys poison today).

Other theorists adopt different strategies here. Asher and Bonevac for example argue that many puzzles in deontic logic are best handled by treating "ought" to be interacting with a non-monotonic conditional. And not by taking wide-scope over it. On their strategy, it is true that:

Adam intends to kill the dean  $\sim\sim$ > Adam ought: (He buys poison today)

and moreover, it is true that Adam intends to kill the dean. However, it is not true that:

X and Adam intends to kill the dean  $\sim\sim$ > Adam ought: (He buys poison today)



One such falsifying X might be that Adam ought to refrain from murder and its means.

I think it's eminently worthwhile to investigate the logical structure of the conditionals we actually use in describing these cases. And intrinsically interesting, as well. However, I observe that Asher and Bonevac's strategy would imply here exactly what Broome doesn't want to say about the akratic. If we analyze what we intuitively express as:

Given that Elena thinks she ought to read her daughter's email,  
she ought then really do so.

as:

Elena thinks she ought to read her daughter's email  $\sim\sim$  Elena ought:  
(She reads her daughter's email).

And if we read the latter as implying, when its antecedent is true, that Elena has at least some defeasible reason to read her daughter's email—as Asher and Bonevac do—then we are saying just what Broome insistently denies about such cases. On his view, the kinds of “oughty” pressures we observe being violated by akratics must be strict and cannot be defeasible. Since I'm not entirely on Broome's side about that, I'm less sure what to think here.

One objection often made against views like Broome's is that if the only relevant oughts that Adam and Elena are subject to is to make a material conditional true, there will be multiple ways for them to discharge their duties. They can bring it about that the consequent is true, or bring it about that the antecedent is false. And sometimes we do think that both of these options has something to commend it. But in other cases the options don't at all seem symmetrical. Schroeder discusses other options which seem even less plausible. He supposes Bradley wants to dance and knows that going to the party is a necessary means to his dancing. Obviously, the normative pressure on Bradley should point towards dancing. (Although this pressure might be only hypothetical, if Bradley's desire to dance is contrary to what he's in a position to justifiably desire.) We might summon up some sympathy for the idea that his abandoning his end is also a permissible response to this pressure. (Especially if the end was unjustified in the first place, and the prospect of going to the party makes its unreasonability more acute.) But it's hard to have much sympathy for the idea that rendering the party no longer a means to dancing, perhaps by arranging for it to be interrupted by the police, is an appropriate response to the same pressure. It's irrelevant whether such acts are appropriate simpliciter. The question is whether they could be dischargings of the duty Bradley's under to go the party, conditional on desiring to dance, and knowing the party to be a necessary means to it. It's hard to see how they could be. Yet these are all ways to make the material conditional that:

Bradley desires to dance and knows going to the party to be a necessary means to it \horseshoe Bradley goes to the party

true.

I appreciate these arguments, though there's much more to be debated before things are settled. However, I want to contest the starting assumption that Broome is committed to the deontic symmetry of different ways of making the material conditional true. What is true is that as a wide-scooper he hasn't provided anything to introduce an asymmetry, That's not the same as saying he's committed to a symmetry. And indeed, as we already observed, Broome thinks that the holding of the "ought" with wide-scope over the material conditional is just *part* of what our intuitive judgments about these cases mean. He's also committed to it's *not* being true that:

Adam in fact intends to kill the dean \horseshoe Adam ought: (He buys poison).

But that leaves open that there may be other interesting normative facts in the vicinity, which do introduce some normative asymmetry between different ways that Adam might see to it that the material conditional is true. Of course we'd like to hear what they are. But that's a different complaint than is usually offered.

My own preferred strategy here is to treat "ought"—or whatever normative operator we work with, it might be a less strict operator—as at the first level of analysis taking a binary form. That is, at the first level of analysis we say:

Adam ought: (He buys poison | [given that] He intends to kill the dean).

And then we go on to debate what we can and cannot establish about the logic of the notion so introduced. It may be that ultimately, the notion can be further analyzed in terms of some non-binary normative operator interacting with some familiar kind of conditional. But I think at the stage of inquiry we're now at, it's premature to say so.

Of course, if we're only in a position to treat:

Adam ought to buy poison, given that he intends to kill the dean.

as the holding of some binary normative relationship between his buying and his killing, then it's unclear how that interacts with the holding of non-binary normative claims, such as the (in this case, false) claim that Adam ought to (or ought to intend to) kill the dean. For we wouldn't then be in a position to assert that the same normative operation is present in both claims. However *it's exactly right that it should be unclear how these claims interact*. Even if the

binary “ought” ultimately reduces to the holding of a non-binary ought towards some conditional, and even if it’s the same non-binary ought as is involved in the claim that Adam ought to kill the dean, we should be unsure exactly what logical relationship these claims have.

That is: Broome and other wide-scopers have been most keen to deny that the “oughts” they’re investigating validate *factual detachment*. This is the inference pattern:

```
Adam ought to buy poison, given that he intends to kill the dean.  
Adam does in fact intend to kill the dean.  
-----  
Adam ought to buy poison.
```

But what about different inference patterns, such as *deontic detachment*:

```
Adam ought to buy poison, given that he intends to kill the dean.  
Adam *ought to* kill the dean.  
-----  
Adam ought to buy poison.
```

In many cases this inference pattern has more intuitive plausibility. But in fact it is and *should* be controversial, as well. Many of Broome’s readers had the impression that he, at least initially, endorsed this inference pattern. I’ve heard this assumed in many informal discussions, and I assumed it myself. It is clear though that in later papers he expresses doubt about deontic detachment too; and some who’ve interacted with Broome assure me those doubts were there all along. So I’m not sure what he intended to communicate about this principle, in his early papers. Other advocates of wide-scope approaches have generally been more causal about this, and seem to have assumed deontic detachment is ok. Much work in deontic logic (work that treats “ought” as a “normal” modal operator) also treats this inference pattern as valid.

However, other work in deontic logic doesn’t assume it is, and as we’ll see, some puzzle cases familiar in the literature raise doubts about it. So it really is a hard substantive question what the logical relationship is between these claims. It doesn’t become easier or less substantive by our deciding that it is a single deontic operation involved in each claim.

Hence, I don’t think we lose anything—but premature commitments—by starting out assuming as little as possible about the internal logical structure of the “ought” claims we’re investigating. This is why I don’t like to say I’m advocating a “wide-scope” view, but rather something more general, of which wide-scope theories are just one possible implementation.

What is common to the more general family of views is that they resist factual detachment. We may or may not want to also endorse Broome’s idea that if:

Given that Elena thinks she ought to read her daughter's email,  
she ought then really do so.  
Elena thinks she ought to read her daughter's email.  
Elena doesn't read her daughter's email.

then Elena has definitely done something wrong—violated some simple, non-conditional “ought.” (That is the most natural way to understand “definitely did something wrong,” though remember as I mentioned, that in his later work Broome has reservations about whether the kind of requirement Elena is violating here is one she has any reason to conform to.) Endorsing this idea would amount to accepting that:

Given that Elena thinks she ought to read her daughter's email,  
she ought then really do so.

*implies* that:

Elena categorically ought not: (Elena thinks she ought to read her daughter's email, but fail to read it).

but it need not be equivalent to the latter. Indeed, it need not be equivalent to *any* categorical ought that implies this one. In particular, it need not be equivalent to, nor even imply:

Elena categorically ought: (Elena thinks she ought to read her daughter's email ---> Elena reads the email).

for any interpretation of the “—>”. Whether any of these further claims are true needs to be hashed out.

I think most of what's most interesting to normative theorists here can be productively debated while making minimal assumptions about the internal logical structure of the judgments we make about cases like Adam's and Elena's. What we're interested in is their external logical structure: what inferences are and aren't valid. But we shouldn't be deciding that in any case based on what we've decided about their internal logical structure. That puts the cart before the horse.

## Section 6

I want to briefly discuss two of the puzzle cases I mentioned before. Both of these raise a presumptive challenge to whether deontic detachment really is a valid inference form.

The first case comes from characters like Frank, that I mentioned in Section 1. If you'll remember, Frank justifiably but falsely believes he ought to be maximally non-interfering with his children. That's what he has most justification, on balance, to believe. And though that may make a difference to what he in fact has most reason to do, I claim it's still possible for Frank, on balance, not to have most reason to act that way. That is, Frank's normative beliefs can be false even when they're justified. Frank, then, is unlike the other characters we've been considering in that his deliberative inputs are categorically justified.

Christian has taught me to associate the puzzle that arises here with a famous discussion of Ewing's—though I have yet to read the original. But in any case, I think of this as a form of Ewing's Puzzle. Please do let me know if I'm misunderstanding the historical dialectic.

As it applies to our present discussion, the issue that presses here is that, we have on the one hand, our intuitive judgment:

Given that Frank thinks he ought to abstain from interfering with his children, he ought then really do so.

This is the same judgment we had about Elena, though Elena's normative belief, unlike Frank's, was unjustified. How are we to understand each of these? Broome understands them as:

Frank ought: (He thinks he ought to abstain from interference \horseshoe He abstains from interference).

where, crucially, Broome thinks the "ought" involved is a strict one, rather than a presumptive or pro tanto one. Remember, Broome thinks that if Frank is akratic and falsifies the embedded conditional, he *definitely will have done something wrong*. The wrongness shouldn't possibly be outweighed by other normative facts about Frank's situation. For our purposes, it's not important whether the "ought" takes wide scope over a conditional, as Broome thinks, or whether we treat it more neutrally as some binary normative relation:

Frank ought: (He abstains from interfering | [given that] He thinks he ought to abstain from interfering).

What I do want to retain, for the moment, from Broome's discussion is the idea that the "ought" we have here is strict.

Now, next, in Frank's case we also have it that his normative belief is justified. That is, it's what he ought to think, given his evidence. Or:

Frank ought: He thinks he ought to abstain from interfering.

Now if deontic detachment is valid for the thoughts we have here, then it should follow that:

Frank ought: He abstains from interfering.

In other words, Frank's normative belief must after all be *true*, not false as we were supposing.

As I said before, some philosophers welcome this result. They acknowledge we're fallible about most things, but they think the normative must be different. Justified beliefs about the normative, they think, must also be true. But I tried to motivate hesitation about thinking that.

If we do want to allow normative beliefs to be possibly false even when justified, then we have two options here.

One option is to deny that the inference pattern here illustrated is valid. That is compossible with other, superficially similar-looking inference patterns in fact being fundamentally different, and *they* still being valid. That is, perhaps *this* kind of alleged example of deontic detachment isn't valid, but others are. That's a possibility. But I assume that if the present inference pattern is invalid, we'll have at least some presumption that the present argument is legitimate instance of deontic detachment, and since it isn't valid, neither is that inference pattern quite generally.

The other option is to abandon Broome's idea that the relationship between Frank's normative belief and his proposed action is *strict*. Perhaps it's just that Frank's normative belief constitutes *some reason* for him to do as he envisages, but that he may in fact have more reasons not to do that. So on balance, he may have most reason to sometimes interfere with his children, and his normative belief is false.

This is the strategy I'm most inclined towards. Note if we go this way, though, we give up on what Broome has emphasized as his central argument for believing in hypothetical (or as he conceives them, wide-scope) normative (or as he conceives them, at least potentially normative) relations. I think that's ok. We have plenty of good reasons to acknowledge these without that argument.

Our second puzzle case comes from a discussion of Chisholm's in the 60s. He supposed that you ought to go visit your grandma. And if or given you will visit, you ought to call ahead and say you will visit. Grandma gets very nervous when guests come announced. This is not a means/ends ought, but it looks at superficially like them. And if we understand it in the same way, then it seems we have:

You ought to visit your grandma

You ought: (You call your grandma and say you will visit | [given that] You will visit)

Now if deontic detachment is valid, it seems to follow that you ought to call your grandma and say you will visit—no matter what else is true. After all, if this inference pattern is logically valid, it's still valid no matter what additional premises are added. (We're not thinking about a kind of non-monotonic validity here.) And I forgot to mention one other thing that is true: you're not going to visit you grandma. Yes, you ought to, but you're not going to.

Are we still so comfortable saying you ought to call and say you're coming? Even though you won't in fact come?

So there seems to be some intuitive pressure against deontic detachment here, too.

However, on further reflection, this case seems problematic for our purposes in the same way that Jackson's Professor was. The "givenness" of your duty to call here seems to derive not from the "oughtiness" of your future visiting, but rather from its *presence*. That is, it doesn't seem like the normative relation between your visiting and your calling is on a par with the relation between Adam's end and his pursuing the means. Rather, it seems like what we should say here is:

You ought to visit your grandma

If you will visit, you ought: (You call your grandma and say you will visit).

It's irrelevant to the second claim whether in fact you ought to visit. It matters only whether you will. And in the case as described you won't.

So we don't have here a counter-example to deontic detachment for the kinds of hypothetical "oughts" that Adam and other cases from Section 1 exemplify. We don't here have an example of any "ought" of the same sort.

One NYU student, Jon Simon, proposed a variation of this case which seems a better test for whether deontic detachment holds for "oughts" like Adam's. (I know there is much existing literature on Chisholm's case. I haven't yet reviewed it, so I don't know if this is discussed elsewhere.)

In Simon's variation, we consider not whether you ought to call Grandma and say you're coming. Instead we consider something you know to be a necessary means to visiting Grandma: You have to start walking through the forest, for that's where she lives. That is, we have:

You ought to visit your grandma

You ought: (You start walking through the forest | [given that] You will visit)

We also have:

In fact you won't visit your grandma.

Deontic detachment for these “oughts” would tell us that despite the latter, though, it still is true that you ought to start walking.

I’m of two minds about this result. On the one hand, yes it does seem right to say you ought to start walking. You ought to start walking, and keep on walking, until you get there. Thereby visiting, as you ought. On the other hand, there also seems to be some awkwardness about saying you ought to start walking. What frame of mind are we envisaging you doing so in? “Here I am starting out, but I’m not going to actually go there. At some point I am definitely going to turn back...”

One response we could have to this is to say:

You ought: (You start walking and visit Grandma)

without its simply being true that:

You ought: (You start walking)

since as a matter of fact you’re not going to finish walking and successfully visit. That is, perhaps  $Ought(X \text{ and } Y)$  does not imply  $Ought(X)$ . If so, this would also amount to an abandonment of deontic detachment, since we’re no longer agreeing that:

$Ought(Y)$   
 $Ought(X \mid [\text{given that}] Y)$

implies  $Ought(X)$ . It only implies  $Ought(X \text{ and } Y)$ .

Another strategy here is to propose a different substitute for deontic detachment. Instead of this being valid:

$Ought(Y)$   
 $Ought(X \mid [\text{given that}] Y)$   
-----  
 $Ought(X)$

it’s proposed, perhaps instead it’s *this* which is valid:

$Ought(Y)$   
 $Ought(X \mid [\text{given that}] Y)$   
 $Y$   
-----  
 $Ought(X)$



and in our examples, crucially, Y is false.

These new inference patterns are not falsified by the examples here considered, and I expect them to be generally harder to find counter-examples to. But I'm not sure that they go in the right direction. I worry, for example, about cases where you ought to Y, and you do or will in fact Y, but not for the reason you ought to. If you ought to visit grandma, because it makes her happy, and you will in fact visit her, in order to murder her, is it still clear whether or not you should be starting out? I feel pulled in different directions.

Also, I'm not so sure it's wrong to say you should be starting out in Simon's example: maybe we don't even need to abandon deontic detachment.

I will propose a different way of thinking about these examples, that helps sort out some of the conflicting reactions I have to them. I hope you find it also helps you think more clearly about the cases.

## Section 7

Ok, my proposal is that we first think of the categorical and hypothetical normative pressures that bear on you as parallel domains. Something like two applications running on your computer simultaneously. This doesn't mean they are autonomous, but at this first stage we think about them separately.

The realm I'm calling categorical I associate with what your evidence or reasons really call for—independently of what practical or epistemic choices you in fact make, and why. That is, I associate this with what we in epistemology call “propositional justification.” (That label is unfortunate for several reasons, but it is now well-established.)

And so understood, it's not clear there is anything wrong with deontic detachment. Your reasons really do support visiting grandma, and also starting out in order to get there. In the Ewing Puzzle case, as I said, I am sympathetic to the idea that Frank's justified normative belief—or rather, the justification he has for it—should give him some reason to avoid interfering with his children. For his normative belief to be false, it only needs to be the case that he also has other, stronger, reasons not to do so.

I don't know. Perhaps deontic detachment isn't right in the end, after all. But understanding categorical normativity in this way, I don't yet see a problem with it.

In the hypothetical domain, on the other hand, I'm not sure what logical relationships obtain. But I tend to think of them as mirroring the categorical side. If Categorically Ought(Y) generates categorical justification to X, then I tend to think that actually Y will generate hypothetical justification to X. But I'm not prepared to identify any general principle of this shape. Instead, I'm only relying on this for heuristic guidance.

I don't know what the logic of hypothetical oughts and justification and undermining is, but I do have definite judgments about its presence, as in the many examples we've considered.

There are many open questions about hypothetical normativity, understood as parallel domain in the way I'm proposing. Are hypothetical normative relations robustly fundamentally there on the ground? Or are they just an ad hoc, local conceptual framework we employ when thinking about specific cases. Maybe there's no globally determinate thing to say about what Adam ought to be doing, even in the hypothetical sense. Some of his actual ends may point him in some directions, others in other directions, and it may be far better for him to give many of these ends up, rather than commit to what reflective equilibrium among them would really require.

Even if hypothetical oughts do turn out to be essentially a local tool for thinking about cases, rather than some fundamental fact objectively there on the ground, I still think the way I'm now in the process of describing for thinking about them will be useful.

Another important question I have to leave open is how insulated the hypothetical normative facts are from the categorical side. When determining what Adam ought, hypothetically, to do or think, is it only relevant what he actually intends and thinks? Does he only need to achieve purely subjective coherence? Or is it relevant also what his evidence and reasons really do, categorically support? Is what he ought to do and think, hypothetically, a joint product of the categorical pressures he's under and his actual commitments—his actual beliefs and intentions? I don't know.

So far, I've left so much open that it may well not be clear anything useful has happened. But now here is the distinctive positive proposal I want to make.

When it comes not to "propositional justification" but to what epistemologists call "doxastic justification"—to which of our beliefs are really epistemically appropriately held, not just called for by your evidence—then both categorical and hypothetical justification matter. I presume the analogous fact also holds in the practical realm.

That is to say, in order for your belief to be appropriately held, it's not enough that you have evidence for it, that you believe in a way that matches the evidence, and that you do so "on the basis of" that evidence. It's also essential that you *lack* hypothetical justification for doing otherwise. For example, your senses may give you justification to believe you see a cup of coffee, absent grounds for thinking otherwise or that your sense are awry. And suppose you in fact have no such grounds. But you do in fact believe, groundlessly, that your senses are unreliable. Despite this, you go ahead and believe there's a cup of coffee there, anyway. On the basis of your senses telling you so. Not because you reconsidered your belief that your senses are unreliable. Instead, in apparent defiance of that belief.

Clearly if you do all this you are messed up. Of course, you were already messed up by having the baseless suspicions you do about your senses. But it seems that in proceeding in the ways I described, you manifested even further epistemic defects. Now it's a stronger claim to say that, therefore, your belief about the coffee cup itself is affected, and fails to be appropriately held. But at least in some cases, I do want to say that. That is, I think at least sometimes, a belief that matches and is held on the basis of what you have categorical justification to believe, can fail to be appropriately held, because of its incoherence with other things you merely believe, or merely are under hypothetical epistemic pressures to believe.

Of course, I'm not saying that your belief about the coffee cup is *worse* than your belief about your unreliable senses. So if one has to go, it may well should be the latter. And even if we hold fixed that the belief about your senses is not going to go, I'm not saying either that between

- (i) incoherently believing the coffee cup is there, or
- ii) contrary to your evidence, but more consistently, refraining from the belief in the coffee cup

(i) is always the worse choice. I'm saying that there's something defective and bad manifest in case (i). In may well be that the available alternatives are worse. I do not feel compelled to think that the normatively least bad of the alternatives available to you must therefore be good or reasonable.

Ok, that's my proposal.

I'll close by mentioning one potential downside. I've come to terms with it, but it may give you pause.

Consider a case of the sort I described, where your evidence categorically points to choice A. But suppose that some of your actual beliefs and intentions hypothetically commit you instead to choice B. Now it seems there may be nothing you can do appropriately, in the sense where "appropriate beliefs" are well-founded, doxastically justified beliefs. For I've just said that, given your hypothetical pressures to B, you cannot appropriately choose A, even if you do so in response to the reasons that really support Aing. Not so long as those hypothetical pressures against Aing are in place. But then, neither can you B, since your evidence doesn't support Bing.

In such cases, should you perhaps give up the actual attitudes that induce those hypothetical pressures? Yes, you do have categorical reasons for doing so, since we're supposing them to be categorically unjustified. That's why the pressures they impose are merely hypothetical. But although you have categorical reasons for giving them up, do you necessarily have available any appropriate, well-founded path for giving them up? That is, could lacking them be not just a statically better place for you to end up, normatively? but actually reachable

by you in a dynamically justified way, given the normative position you're now in?

I don't see any reason to think you must. You might luck out, and have the unjustifiedness of these attitudes suddenly just dawn on you. But that's a development you'd be passive with respect to. If it happens, great, but it's not the kind of thing you're going to dynamically rationally make happen. Absent any such sudden realization, I see no reason to think a dynamically rational path from the attitudes you now have to the attitudes your evidence and reasons in fact call for must always be available. Some of the unjustified attitudes you have, that don't cohere with the evidence, might always obstruct you from walking such a path in a way that's in fact well-founded or dynamically rational.

In other words, your situation may well be a normative tragedy. You shouldn't be where you are, but none of the paths open to you can be dynamically rationally walked, either.

I've given no constructive proof this will be so, but I see no reason to expect it won't be.

And recall what I assumed to generate this structure: merely that some of your actual attitudes were contrary to evidence, and so imposed hypothetical normative pressures on you that were also contrary to evidence. Being the imperfect creatures we are, this kind of scenario is ubiquitous.

Hence a threat of my view is that it makes normative tragedy, of both epistemic and practical sorts, nearly as ubiquitous.

As I've said, I've come to terms with this and started to believe it's true. I guess it's ironic that someone who's argued (in terms of propositional justification) for anti-skepticism should come to think (in terms of doxastic justification) that not only do our beliefs tend to be unreasonable, but that they are moreover inescapably and tragically so. Perhaps I'm over-reacting.

As it happens, I think we are lucky in that, often when we are in positions where our beliefs and intentions fail to fit the evidence, we sometimes do "jump" to overall positions that are more conformant. Sometimes we just listen to some normative pressures and blindly ignore others. And sometimes this makes us better off, and in the end more reasonable. Perhaps you are more reasonable when you believe the coffee cup is there than if you more consistently heed to your baseless doubts in your senses. We are lucky that this is so. If my proposal is right, this is a kind of "luck" because the jump that was involved, though it went from one set of attitudes to another that better fit your reasons, was nonetheless not dynamically rational.