

# EPISTEMIC DEONTOLOGISM AND ROLE-UGHTS

Jon ALTSCHUL

ABSTRACT: William Alston's argument against epistemological deontology rests upon two key premises: first, that we lack a suitable amount of voluntary control with respect to our beliefs, and, second, the principle that "ought" implies "can." While several responses to Alston have concerned rejecting either of these two premises, I argue that even on the assumption that both premises are true, there is room to be made for deontology in epistemology. I begin by offering a criticism of Richard Feldman's invaluable work on 'role-oughts,' whereupon I develop my own positive view in light of Feldman's shortcomings. The upshot is that while we as epistemic agents are not responsible for the beliefs we form, we are nonetheless responsible for the various bodily or mental activities that typically bear a causal influence on belief formation.

KEYWORDS: justification, doxastic voluntarism, role-oughts, epistemic responsibility

## 1. Introduction

Most of us think that there are various things that we as moral agents ought to do, like picking up your litter off the ground or calling a close friend when she needs a shoulder to lean on. When we fail to do the things we ought to do, or conversely we do the things we ought not to do, we are rendered morally accountable for how we acted, and blame may be in order. It is less clear, however, whether a similar mode of assessment can be applied, not to the actions we perform when interacting with the world, but to the beliefs we form in light of our epistemic situations. Is it true that there are certain beliefs we as epistemic agents ought to hold, and moreover does failing to believe as we ought (or believing as we ought not) leave us epistemically responsible for our doxastic states?

Those who answer this question in the affirmative tend to conceive of epistemic justification in what William Alston calls a "deontological" way, as having to do with obligation, permission, requirement, blame, and the like.<sup>1</sup> This way of thinking about justification would be especially appealing to the philosopher who holds that at the crux of the concept of justification is the idea of being receptive to the reasons one has and using them to shape one's doxastic

---

<sup>1</sup> William Alston, "The Deontological Conception of Epistemic Justification," in his *Epistemic Justification: Essays in the Theory of Knowledge* (Ithaca: Cornell University Press, 1989), 115.

livelihood. Justification, understood deontologically, is the mark of an *achievement* on the agent's part – to believe as one ought implies that the agent has fared well or has done a good job in the face of her epistemic situation. This is a feature that has traditionally been underemphasized or omitted altogether by various reliabilist or other externalist theories. If what matters most about justification to an externalist is that there be in place some suitable connection between the agent and the world, then the very idea of what the *agent* is responsible for seems to go out the window.<sup>2</sup> We wish not to surrender either of these highly plausible features about justification: 1) justification places various standards upon the epistemic sturdiness of the specific belief formed – that is, its connection to truth, but 2) it also places various demands on *how* the agent should arrive at the belief. Given the strength of an agent's evidence and cognitive faculties, there are various things that *she*, in that moment, is *supposed to do*.<sup>3</sup>

There is an important and much discussed problem for this deontological way of thinking about justification, which comes from Alston in an oft-quoted passage on the subject:

Now this conception of epistemic justification is viable only if beliefs are sufficiently under voluntary control to render such concepts as *requirement*, *permission*, *obligation*, *reproach*, and *blame* applicable to them. By the time-honored principle that 'Ought implies can,' one can be obliged to do A only if one has an effective choice as to whether to do A.<sup>4</sup>

Belief seems not to be under the influence of the will in the same way as are our ordinary actions, like wiggling a toe or paying the rent. Borrowing Alston's own example, when I see a car coming down the street, I am suddenly struck with the belief that there is a car on the street. There is no real sense in which forming

---

<sup>2</sup> See Laurence Bonjour, *The Structure of Empirical Knowledge* (Cambridge, MA: Harvard University Press, 1985) for a classic objection to externalism on this score. For some emended versions of reliabilism that purport to avoid this objection, see Alvin Goldman, *Epistemology and Cognition* (Cambridge, MA: Harvard University Press, 1986) and John Greco, *Putting Skeptics in Their Place* (Cambridge: Cambridge University Press, 2000).

<sup>3</sup> The targets of Alston's original argument were those philosophers who sought to *analyze* the concept of justification from an entirely deontological point of view (e.g. Carl Ginet, *Knowledge, Perception, and Memory* (Dordrecht: D. Reidel, 1975) and Roderick Chisholm, *Theory of Knowledge*, 2<sup>nd</sup> Edition (Englewood Cliffs, NJ: Prentice-Hall, 1977)). I wish not to defend this very strong version of deontologism. Instead, all I am interested in investigating is whether there exists a deontological dimension or feature to the concept of justification, leaving it open whether there are other features relevant to satisfying this concept. Note however that Alston's argument, if successful, would defeat even this weaker form of deontologism.

<sup>4</sup> Alston, "The Deontological Conception," 118.

this belief was done out of my own willing or choosing. As others have pointed out along with Alston, we are simply at the mercy of our evidence and reasons when arriving at a doxastic position.<sup>5</sup> For this reason, it is hard to see how deontological concepts could have any application in epistemology.

The purpose of this paper is to vindicate epistemic deontology in the face of Alston's *argument from doxastic involuntarism*. Other recent attempts to respond to the argument have commonly challenged either of Alston's two key premises: the claim that we lack a suitable amount of voluntary doxastic control and the principle that "ought" implies "can." Unlike these responses, I do not wish to challenge Alston's assumptions in this argument. I believe they are both correct, but I shall not argue for this in this paper. I am instead interested in answering whether deontology can have any place in epistemology in light of the truth of such assumptions. Ultimately I will claim the answer is "yes," and I will do so primarily by presenting a criticism of Richard Feldman's invaluable work on role-oughts – according to which there are duties and obligations that get assigned to agents solely in virtue of occupying some role – which he has presented as a way of answering Alston's argument. While, as I will argue, Feldman's answer to Alston's argument ultimately fails, his approach provides important insights to what the nature of our epistemic obligations are. Specifically I will argue, mirroring a position from John Heil,<sup>6</sup> that while we are not responsible for the beliefs we form, we are responsible for the bodily and mental actions we take that cause us to form our beliefs. These actions are ones over which we typically do have voluntary control and moreover are the sorts of actions that are characteristic of responsible epistemic agency.

## 2. Role-Oughts and the "Ought implies Can" Principle

Alston's argument claims that having a suitable amount of voluntary control is a necessary condition for attributing a deontological judgment toward one's belief. There are two central strategies for responding to Alston's premise. The first is to accept his premise as true but hold that we do possess the requisite form of control

---

<sup>5</sup> John Heil, "Doxastic Agency," *Philosophical Studies* 43 (1983): 357, Richard Feldman, "Voluntary Belief and Epistemic Evaluation," in *Knowledge, Truth, and Duty*, ed. M. Steup (Oxford: Oxford University Press, 2001), 83, Nikolaj Nottelmann, "The Analogy Argument for Doxastic Voluntarism," *Philosophical Studies* 131 (2006): 560-61.

<sup>6</sup> Heil, "Doxastic Agency."

with respect to our beliefs.<sup>7</sup> The second strategy concedes to Alston that we lack such control but denies that it is necessary for deontological assessment.

For our purposes, we will investigate the latter strategy only, and we will start with one its chief proponents, Richard Feldman.<sup>8</sup> According to Feldman, the first strategy will not succeed because the kind of control necessary to satisfy Alston's premise is what Feldman calls *response control*. One has response control with respect to  $\Phi$ ing when "one has the ability to use practical reason to respond directly to incentives and the like."<sup>9</sup> With our more ordinary actions, we typically do have such control. Suppose you ask me to pass the salad dressing to you across the table. Here the fact that you are in some need and I wish to be a good friend is a practical reason I now have for performing this action. And I can right now in that very moment, directly on the basis of this reason, either pass the dressing over to you or do something else, like say "No" or agree but only if you pass me the roles. In contrast, if instead of passing the dressing, you asked me to believe right now that Oprah is the Canadian Prime Minister, this is something I cannot do. You might even offer me some kind of monetary reward for complying, yet no matter how big the incentive may be, there is nothing I can do at this moment directly on the basis of that incentive to form the requested belief. It is not a matter of how much I want you to offer me; rather, it is a matter of what I am and am not able to do at this moment.

If we lack response control with respect to our typical beliefs, does it follow that there is no basis to epistemically evaluate such beliefs from a deontological standpoint? Not at all, says Feldman. Instead, he thinks we can provide such evaluations with regards to how well the belief formed is a proper response to the information presented to the agent at the time. Consider the deontological term,

---

<sup>7</sup> For defenders of this first strategy, see Matthias Steup, "Doxastic Voluntarism and Epistemic Deontology," *Acta Analytica* 15 (2000) and "Doxastic Freedom," *Synthese* 161 (2008), Sharon Ryan, "Doxastic Compatibilism and the Ethics of Belief," *Philosophical Studies* 114 (2003).

<sup>8</sup> Richard Feldman, "Voluntary Belief and Epistemic Evaluation," and "Modest Deontologism," *Synthese* 161 (2008). For other defenders of this second strategy, see Jonathan Adler, *Belief's Own Ethics* (Cambridge: MIT Press, 2002), Philippe Chuard and Nicholas Southwood, "Epistemic Norms without Voluntary Control," *Nous* 43 (2009), and Conor McHugh, "Epistemic Deontology and Voluntariness," *Erkenntnis* 77 (2012).

<sup>9</sup> Feldman, "Modest Deontologism," 347. Feldman's notion of response control mirrors the kind of voluntary control espoused by Jonathan Bennett, "Why is Belief Involuntary?" *Analysis* 50 (1990), who conceives of voluntary control as being responsive to practical reasons. See Pamela Hieronymi, "Controlling Attitudes," *Pacific Philosophical Quarterly* 87 (2006) for a discussion of Bennett and also her own notion of what she calls managerial control. For an understanding of voluntary control similar to the one being described by Feldman, see McHugh, "Epistemic Deontology and Voluntariness."

“ought.” It is reasonable that if any deontological terms are applicable to belief, we should expect at a minimum that statements of the form,

(OUG) You ought to believe that p,

can be, and oftentimes are, true.<sup>10</sup> Feldman claims that OUG-statements can be true, and thus there is room to be made for at least a modest version of epistemic deontology.<sup>11</sup> His argument runs as follows: OUG is equivalent to:

(RES) Believing that p is the epistemically appropriate response to S’s evidence.

Since RES-statements can be true, says Feldman, the same goes for OUG-statements. Suppose your current evidence consists of a visual experience of a computer screen ahead. Assuming you have no reason to think you have taken some mind-altering drug or are looking at a hologram, believing that there is a computer screen ahead is the appropriate response to the evidence you possess. Thus, the relevant RES-statement is true in this case. But if forming this belief in response to the evidence in this situation is the one that is appropriate, then it must be true that you ought to believe that there is a computer screen ahead.

The merits of Feldman’s argument is that we find a way to make room for at least some kind of deontological assessment of agents’ beliefs, at least with respect to what they ought or ought not to believe.<sup>12</sup> Also on Feldman’s account, the precise sort of deontological assessment the agent receives is completely divorced

---

<sup>10</sup> To elaborate I will assume for this paper that whenever it is true that S is obligated, required, or has a duty to believe that p, then it will also true that S ought to believe that p; and also whenever it is true that S is permitted or cannot be blamed for believing that p, then it will not be the case S ought not to believe that p. Thus, it will turn out that no deontological statements about one’s belief can be true unless there is some corresponding ought-statement that is also true.

<sup>11</sup> There is one sense in which OUG-statements clearly are true, as when one makes a claim about the normal behavior or expectation of the discussed agent (see Nicholas Wolterstorff, “Obligations of Belief: Two Concepts,” in *The Philosophy of Roderick Chisholm*, ed. E. H. Lewis (LaSalle, IL: Open Court, 1997)). Knowing that their child is afraid of her room in the dark, one parent might correctly say to the other ten minutes after putting the child to bed, “She ought to be thinking there are monsters in her closet by now.” The relevant ‘ought’ here does not, nor is it meant to, describe whether the child would be in a good epistemic situation by forming this belief, but only whether this is the doxastic state the parents would expect the child to be in at a certain point. Thus we would not want epistemic deontology to rest solely on this separate class of ought-statements. I will henceforward disregard this separate class for our discussion.

<sup>12</sup> Feldman’s deontology is called *modest* in the sense that he thinks a minimal number of deontological concepts may be applied to belief – specifically the concept ‘ought’ and ‘ought not’ – without being thereby committed to the applicability of other deontological concepts, like ‘duty,’ ‘permission,’ or ‘requirement.’ See Feldman, “Modest Deontology,” 348.

from the question of whether the agent exercised response control with respect to her belief. We can assume that once you look at the computer screen before you, you are, in a sense, “locked” into forming that corresponding belief. There is nothing you can do at this moment to believe otherwise, no matter how much of an incentive I may offer you. Nevertheless, says Feldman, there does exist a proposition which, given your evidence, is the one that you truly *ought* to believe.

Despite these merits, I will argue that Feldman’s argument fails. For, there are plenty of real life situations in which despite the fact that believing *p* would be the appropriate response to the agent’s evidence, the agent is simply unable to believe it. Perhaps a mother whose son has been missing for weeks has been told by the competent lead detective that the son was found dead, yet her overwhelming hope that he remains alive makes it psychologically impossible for her to accept what the evidence actually supports. If OUG were equivalent to RES, then it would follow that there are some propositions one ought to believe but which one cannot; yet this would stand in violation of “that time-honored principle that ‘ought’ implies ‘can.’”<sup>13</sup>

Feldman is unbothered by this objection because he rejects the “ought implies can” principle outright.<sup>14</sup> He does so by introducing the notion of ‘role-oughts,’ according to which “[t]here are oughts that result from one’s playing a certain role or having a certain position.”<sup>15</sup> As an apartment renter, I *ought* to pay the rent by the first of the month; this is something that I am obligated to do. Surely I do have the relevant kind of control over whether I fulfill or violate this obligation (it is within my power to purchase a brand new television instead of mailing in the rent check); but the fact that I have such control seems not to be the only, or even the primary, factor as to whether such an obligation to pay the rent exists. For, even if I were unable to pay the rent in a given month (say, I was fired from my job, and I now have insufficient funds), this seems not to undermine the fact that I still ought to pay the rent. My obligation to pay the rent is binding whether or not it is within my control to pay. According to Feldman, the source of the obligation – the fact that I ought to make this payment – rests on the fact that I have taken on the role of a renter.

---

<sup>13</sup> Alston, “The Deontological Conception,” 118.

<sup>14</sup> Others who have rejected this principle include Michael Stocker, “‘Ought’ and ‘Can,’” *Australasian Journal of Philosophy* 49 (1971), Walter Sinnott-Armstrong, “‘Ought’ Conversationally Implies ‘Can,’” *Philosophical Review* 93 (1984), and Ryan, “Doxastic Compatibilism.” Chuard and Southwood (“Epistemic Norms without Voluntary Control”), rather than claiming that “ought implies can” is false, argue that it cannot be used effectively to reject epistemic deontologism.

<sup>15</sup> Feldman, “Voluntary Belief,” 87.

If we could extend this notion of a role-ought to the case of occupying the role of a believer, then the question of whether OUG can be true would seem to not depend on whether forming that belief is something the agent can or cannot do. However, this strategy will not work. For, it is evident that there is a substantial disanalogy between taking on the role of a renter and taking on the role of a believer. In the former case, the obligation is contractual. Even if we allow that my obligation to pay the rent is explained by the fact that I am a renter, it is clear that it was of my own choosing to take on this role. I agreed to become a renter when I signed the landlord's lease.<sup>16</sup> I brought this obligation on myself. Furthermore, had I been forced beyond my will to sign the contract (due, say, to some compelling coercive power) it is not so obvious that I would still be obligated to pay the rent. At the very least I could give a good argument to the judge as to why I felt I did not owe the landlord any money. But, in the case of the role of a believer, no such contract was ever presented. Indeed, it is not even possible for any sort of contract to be presented prior to one's agreeing to become a believer. Agreeing to a contract of any form requires that *one already have certain beliefs* about what the contract says and what one has to do in order to agree. We are believers, a fact over which we have no choice at all.

Feldman is aware of this disanalogy, and his answer is to present other examples of roles, occupations of which generate obligations, that more closely align with our occupying the roles of believers. He says:

Teachers ought to explain things clearly. Parents ought to take care of their kids. Cyclists ought to move in various ways. Incompetent teachers, incapable parents, and untrained cyclists may be unable to do what they ought to do.<sup>17</sup>

If a parent, for example, is so overcome by her alcoholism that she is unable to take care of her kids, that seems not to take anything away from its being the case that she ought to be doing this. In response to this suggestion, it is apparent to me that the roles Feldman mentions in the above passage are still roles that people choose to occupy. Teachers choose to enter the profession, a couple chooses to have a child, and even cyclists make that initial decision to start riding a bike. That being the case, there is no reason here to suppose that the obligations associated with these roles are not contractual in nature. Consider a similar question to the one regarding a renter forced into signing a lease. If one had no choice over whether to occupy one of these roles, would the obligation associated with that role still exist? When a woman becomes a mother as the result of a rape,

---

<sup>16</sup> See Richard Feldman ("Epistemic Obligations," *Philosophical Perspectives* 2 (1988)) for a discussion of contractual obligations.

<sup>17</sup> Feldman, "Voluntary Belief," 87-88.

Jon Altschul

and she has no inclination or desire to be a parent, is it still the case that she ought to take care of her kids? As with the coerced renter, it is not so *obvious* that the victimized mother remains under an obligation, given that her occupation of this role was forced upon her. Given that the obligations associated with being a believer are in no way contractual, and could not possibly be, I think we should agree that such obligations must be disassociated from Feldman's notion of role-oughts.

However, since his initial discussion, Feldman has given us a response to this newer objection:

One can argue that one can only have role oughts with respect to voluntarily adopted roles. I do not see why this is true. There are ways one ought to eat – chew before swallowing – but the role of eater is not a role one took on voluntarily. There are ways one ought to breathe, even if this is not a voluntarily adopted role.<sup>18</sup>

Feldman is correct that being an eater and being a breather are roles that, like being a believer, no one adopted voluntarily; and in this respect they serve as better analogies than the ones discussed above. But we face a different sort of disanalogy here. Insofar as Feldman is granting that: 1) taking on the role of a believer is involuntary and 2) we have no control (response control) over whether to believe, say, that *p*, then we should expect the right kind of analogy to be one where not only is the adoption of the role involuntary, but also *that the activities associated with that specific role are ones over which we have no control either*. So, while it is true that we lack control over whether to *be* eaters or breathers in the first place, as with being a believer, we nonetheless typically display a considerable amount of control over *how* we execute our tasks in these roles. Granted, we do commonly eat and breathe with hardly any conscious attention to these activities as we do them. But regardless of how little or much I may concentrate on my own chewing habits as I eat, at any given moment whether I chew on my food more or I cease chewing and swallow is something completely up to me – as much as is wiggling my big toe. Consider a friend who offers you \$10 to stop chewing on your steak now and swallow that bite whole. It may give you a considerable stomachache to comply, but surely it is up to you to either accept or reject the offer. That being the case, Feldman is ill-advised to use these sorts of roles to establish the falsity of the “ought implies can” principle. Let me elaborate.

What are we to say about the individual who, due to some rare abnormality, never developed any teeth and is all gums from ear to ear? When a

---

<sup>18</sup> Feldman, “Modest Deontology,” 351.



plate of filet mignon is placed in front of (let us call him) Toothless Mortimer at a wedding reception, we may ask, “As an eater, what ought Mortimer to do?” Is it true of Mortimer that he ought to chew his food? I do not see how Feldman can say otherwise. Given that Mortimer has taken on the role of someone who eats, that one who eats ought to do so properly, and that proper eating involves chewing before swallowing, it seems that Toothless Mortimer ought to chew his food. But, surely this is an odd result. As Mortimer looks down at his steak and asks, “What I am supposed to do here?” consider how insulting and demeaning it would be for another diner at the table to tell him to start chewing his food.

If I am right that it is false that Toothless Mortimer ought to chew his food first, this shows that Feldman’s eating example fails to demonstrate situations in which one ought to do something that one cannot. Of course, Feldman might respond by suggesting that when he said “[t]here are ways one ought to eat,” he was speaking more broadly than I have been presuming with the above example. The idea would be that, in general, people ought to chew before they swallow, but also that this rule is not meant to apply across the board. Perhaps Feldman would allow for exceptions, as might be the case with Toothless Mortimer. Given that he has no teeth, of course it would be silly to say that he ought to chew. Fair enough, but the question then becomes: *why* do we allow exceptions for people like Toothless Mortimer but not for other people? It seems the only answer available here is that Toothless Mortimer, given that he has no teeth, is simply *not able* to chew his food. That is to say, he cannot chew his food. But if we let this explain why exceptions to this norm are allowed, what this boils down to is that it is not the case that Toothless Mortimer ought to chew his food, *for the reason that he cannot*. In other words, it would be the “ought implies can” principle explaining why Toothless Mortimer is relieved of his obligation. But this was the very principle Feldman was attempting to reject. Hence, this response will not succeed.

To generalize the above point, when considering the statement, “S ought to  $\phi$  in virtue of occupying role R,” I find it highly implausible that a statement of such a form will hold in all situations without exception to all those who take on the role of R. Even if we were to try to craft the relevant  $\phi$  associated with the role of eating such that even subjects like Toothless Mortimer would now fall under the relevant obligation (e.g. “eaters ought to eat with attention and not spill any food from their mouths”), there are some eaters out there who would still be relieved of the associated obligations. (Consider a completely immobilized quadriplegic who has lost all motor functioning in her hands and face and must be hand-fed by another. Such a person qualifies as an eater as much as does anyone else, yet it is untrue that she ought to eat attentively and not spill.) And when we

ask why such an exception is allowable, I do not see any tenable way to answer this question without appealing to certain capabilities that the subject lacks. So, aside from whether “ought implies can” is a true principle or not, my conclusion is that Feldman’s use of ‘role-oughts’ is inadequate to reject the principle.

### 3. Sleeping, Tasting and Seeing

Based on the considerations raised in the preceding discussion of Feldman’s role-oughts, we have seen that if there is such a thing as epistemic deontologism at all, and furthermore we wish as Feldman does to model the obligations associated with our beliefs on some sort of activity or role that we normally take to be open to deontological assessment, then this activity or role must be one that we not only occupy non-voluntarily, but also one over which we lack any response control with respect to that activity. In this section I turn to offering two such roles that fit this description; thus they serve as better analogies to the role of being a believer than the alternatives Feldman has given us. This discussion will then serve as a basis for helping us to answer whether there are any such things as epistemic obligations, and if so, what the nature of those obligations are.

Before we proceed, however, let me note that one benefit of the view I will propose is that, unlike for Feldman, we will not be forced to reject the “ought-implies-can” principle. All I have said so far is that Feldman has not demonstrated with his role-oughts that the principle is false. I believe “ought implies can” is a true principle, but I have not defended its truth, nor do I intend to do so here. However, Alston is surely right when he describes that Kantian Dictum<sup>19</sup> as one that is “time-honored,” and although this is no adequate justification for its truth, it at least moves us to question whether there can be any room for a concept of epistemic deontology consistent with its being true. This is the question I wish to answer now. So, I will henceforward be *assuming* that “ought” does imply “can.”

One role that aligns more analogously with being a believer would be the role we each have of sleeping. It is important that we each get a certain number of hours per night, without which there can be harmful consequences. As with believing, being a sleeper is a role none of us adopted voluntarily. But also like believing, it turns out that we have very little control when it comes *to* sleeping. Consider as the night gets later, and a parent walks by little Timmy, who is playing with his toys in his bed, and says, “Hey Son, you ought to sleep now.” If we take this claim literally, the thing that Timmy is supposed to do right now is sleep. But he cannot do this, not “just like that.” Granted, if he lies back in his bed

---

<sup>19</sup> Immanuel Kant, *Critique of Pure Reason* (London: Macmillan & Co., 1933): 473 & 637.

and closes his eyes, these are bodily actions that typically *precede* being asleep. But sleeping and closing one's eyes are not the same. Once in bed with the lights out and eyes closed, and after he has counted all his sheep, there is nothing left for Timmy to *do* so as to be sleeping. The state of sleeping is something that just happens to one with little active participation by the agent. In this sense, I think Timmy, as well as most of us, has no control over whether he sleeps right now or not.<sup>20</sup>

My suggestion here is that the sense in which we lack control with respect to sleeping serves as a better analogy to the role we take on as believers than eaters or breathers. What, then, are the obligations we have in virtue of being sleepers? I propose that one obligation we do *not* have is the obligation to sleep, if by "sleep" we mean to be in the state of being asleep. Usually, when we hear someone say to us, "You ought to sleep" (or "You ought to get some sleep"), we know exactly what instructions are being handed down. The instructions are not to be asleep, even if, as matter of social convention, this is something we are often *told* to do (said to Timmy: "You had better be sleeping by the time I come back to check on you!"). Rather the instructions are only to get into bed, turn out the lights, close our eyes and then wait for sleep to wash over us. Is there really anything else besides performing these actions that Timmy can be expected to do so as to comply with carrying out his parent's instructions? Suppose Timmy stubbornly retorts to his mother, "What do you mean, 'go to sleep,' Mom? What am I supposed to *do*?" If the mother simply said, "Look, I don't know how you do it. Just sleep!" there really isn't much more for Timmy to do. There is no sort of willing that Timmy could do in that moment to do what his mother just told him. (Consider, in contrast, if the mother told Timmy to wiggle his finger, and he asks how, here I think it would appropriate and correct for her to say, "I don't know *how* you wiggle your finger. Just do it!")

If after turning out the lights and closing his eyes, Timmy then has a restless night where he just does not ever fall asleep, would it be fair of his mother to say the next morning: "You didn't sleep last night? But I told you to right before

---

<sup>20</sup> Certainly there is no basic control in this case. Sleeping is not like moving your finger. But there is no non-basic control either. Timmy can control his body – e.g. be in bed or be in the TV room, have his eyes opened or closed – but it is wrong to think that these activities are *part of* the action of falling asleep, as cutting the vegetables is a part of the action of making dinner or getting out a checkbook is part of the action of paying the rent. Unlike cutting the vegetables or getting a checkbook, moving one's body to bed and closing one's eyes are activities that come before, and usually facilitate, the further event of being asleep. Cutting vegetables, while the agent tends to do this at the beginning of preparing a meal, is itself *a part of* the (complex) activity of making dinner.

bedtime! Were you not listening to me?” It would not. There are no grounds for the mother to hold Timmy accountable for his sleeping or lack thereof, because although there are a set of actions that typically causally precede the event of sleeping (e.g. turning out the lights), they do not always result in sleep; and when they fail to terminate with this result this can be through no fault of the agent that the intended performance fell short. Thus, I contend that while there do exist obligations associated with our roles as sleepers, the obligations we have involve doing various actions that tend to causally precede sleeping, and they involve nothing regarding *being* in that intended state.

The second example I wish to discuss has to do with the kinds of experiences or sensations we have. Let me first concede that unlike eating or believing or sleeping, the very notion of having obligations in virtue of occupying the role of a perceiver seems to be somehow misplaced. Perception occurs at the sub-personal level, after all, and we have very little active participation in what we perceive. Thus with regard to evaluating an experience one came to have on a given occasion, the evaluation is more suitably directed at the perceptual system of the agent, and not the agent herself. Nevertheless, there are occasions where we do hold agents accountable for what they see or hear or feel. We say things to each other like, “Oh, you ought to taste my ice cream flavor!” The “ought” here is merely prudential; it is not a requirement of what you categorically are supposed to be tasting now. Even still, assuming you wish to comply with what your friend has requested, what is it that you are supposed to do? I claim that it is untrue that you are supposed to taste the ice cream, where by this I mean *have* a certain gustatory experience. Rather, as with the case of sleeping, what you are supposed to do is perform various bodily actions that tend to causally result in the intended state being produced. So, to do as your friend has instructed, the thing to do is take hold of the ice cream cone and then lick the treat with your tongue and lips. There is no further requirement after these motions in which you are to *then have* the gustatory experience. The experience is what simply happens once you have done the necessary bodily actions. Even if it does not – if, say, you lick the ice cream but then no gustatory experience follows, perhaps because your taste buds have suddenly been anesthetized – it is through no fault of your own that you did not do as your friend asked. The fault lies instead with the gustatory system.

In a similar fashion we are sometimes, though not commonly, handed obligations about what to see or hear. Suppose two concerned parents have allowed their teenaged daughter, Ellie, to go out with her friends to see the new horror flick, but – fearing her exposure to all the gratuitous violence – they agree only on the condition that Ellie not see any of the violent parts. At first glance, we

may view this as an unreasonable demand on Ellie for the simple reason that none of us, in general, has any control over what we perceive. As my eyes are gazed upon this computer screen right now, I cannot help but have an experience of a computer screen. In a similar way, when I place my hand on the hot stove, there is no sense in which I have control over whether I feel pain or not; the pain is something that just washes over me.

Even though Ellie cannot directly control what experiences she has at any given moment, she nevertheless does have considerable control over whether she complies with her parents' demands. For she can willfully perform various actions that tend to result in her not having the selected experience. As the scene of the cheerleader walking up to the attic all by herself begins, and the darkened, creepy music elevates, Ellie knows that something bloody is about to happen; it is up to her in that moment to, say, put her head down to look at her lap, place her hands over her eyeballs, excuse herself from the group and go to the restroom, or some similar sort of action. In doing so, Ellie winds up seeing none of the gore on the screen, and I argue she has fulfilled her responsibility to her parents' demands.

Just like with the case of sleeping, parents may *say* to kids like Ellie, "You ought not see any of the violent scenes of the movie," but the specific obligation being delivered is not an obligation about what seeings or visual experiences they are supposed to have. The obligation is instead to perform a set of actions that the agent has reason to think tends to result in the omission of the intended experience. As with the ice cream example discussed above, we can imagine that after Ellie reaches the lobby for the purpose of avoiding her seeing the violent scene, the movie theater – because of some new marketing promotion – has placed miniature screens all over the lobby and restrooms of the same horror movie running at the exact same time. When Ellie reaches the lobby, she is caused to see the cheerleader get slashed by the killer. Has Ellie violated her obligation to her parents here? Although she saw the death scene her parents did not want her to see, consider how unreasonable it would be for them to punish Ellie in this situation. Most people are like Ellie in that they assume the lobby is a safe place to hideaway from a running movie. There is nothing more that could be expected of Ellie to have thought of or done so as to achieve what her parents wanted. While the result is unfortunate in the case, the unfortunate result is due to circumstances beyond anything for which Ellie could be responsible.<sup>21</sup>

---

<sup>21</sup> When the obligations are directed at having certain experiences, and when they are directed specifically at doing certain bodily motions associated with our sense organs, certainly come can come apart. A parent who is about to be executed may say to her child, "Look away." This is quite different from the situation where the parent, noticing her child staring directly into the

#### 4. Responsible Epistemic Agency

Alston holds that the “ought implies can” principle underwrites his key premise that there can be obligations, duties, permissions, etc., about what one ought to believe only if one has a suitable amount of voluntary control over what one believes. While some have questioned the truth of this premise, and others have challenged Alston’s further premise that we lack the requisite form of doxastic control, I believe that Alston is correct here on both accounts. Thus, I am in agreement with Alston that while we do often say them to each other, claims like “You ought to/have a duty to/are required to believe that p” are never strictly (epistemically) speaking true. However, does it follow from this that there is no room for deontological epistemic evaluations when it comes to an agent’s doxastic performance? I argue that it does not. Even though we may grant that there are no requirements about *what* one is supposed to believe in a given situation, we can still concede the plausible idea that there is a deontological dimension of epistemic justification.

How is it, then, that epistemic agents can be assigned responsibilities with respect to their doxastic performances, and yet it is untrue that there are certain propositions one ought (or ought not) to believe? The answer I wish to motivate in this final section echoes a position espoused by John Heil who says:

It is not that one has a choice in the beliefs that one forms, but that one has a say in the procedure one undertakes that leads to their formation. The notion of ‘epistemic responsibility’ attaches to the undertaking of appropriate procedures.<sup>22</sup>

Heil’s point aligns nicely with the non-epistemic examples discussed in the previous section. Timmy has no choice about whether he is in a state of being asleep or not, but he does have a say in the procedure he undertakes that typically leads to sleep. Ellie has no choice about what visual experiences she has, but she does have a say in the actions she performs that leads to her not seeing the blood and gore. The choices these children make with regard to those actions that lead to the intended state are sufficient for us to evaluate whether the child has or has not done what she ought. Yet, this evaluation is directed not at the sleeping state or the experience the child winds up having, but at the actions taken that causally precede these states.

---

Sun, says “Look away.” The former is a case in which the parent wants the child not to have the emotionally damaging experience of seeing his parent die. In the latter case, parents say this sort of thing to their children not because of the visual experience they would have by seeing the bright Sun (though this can be an unpleasant experience), but rather because of the damage the Sun’s rays can cause to the child’s retina.

<sup>22</sup> Heil, “Doxastic Agency,” 363.

Doxastic evaluations, I argue, work in the same manner. While we do not choose which propositions to believe, we have a great deal of choice in the actions we perform that lead to a belief's formation. Some of these actions may be bodily actions, and we can perform these responsibly or irresponsibly. If one is in a state of doubt that one's experience does not match one's outer environment, it may be one's epistemic obligation to move one's body around to get a better look or to alter the lighting conditions to make the subject matter more clear. Other sorts of bodily investigations may be necessary for beliefs sourced in, say, memory or testimony. (e.g. one might be required in some situations to verify through acquaintances that the testifier is a reliable person.) Thus, the agent may be epistemically required to *do* various things so as to acquire some new evidence on the matter. Once it has been acquired, however, it is the evidence that then determines what belief will be formed, not the agent herself.

In cases of purely critical or deliberative reasoning, when there is no need at this point to acquire additional, independent evidence, there are still certain mental actions the agent can perform that are relevant to the formation of a belief and that are sufficiently under one's control. Suppose Sam is nearing the moment of forming an opinion as to whether *p* is true or not. She is in possession of both a set of epistemically relevant reasons concerning *p* (labeled R1, R2, and R3) and also a set of epistemically irrelevant mental states that nonetheless tend to have causal influence over people's beliefs (labeled M1, M2, and M3). M1 might be, for instance, Sam's deep wish that *p* be true, while M2 is a feeling of hatred she has for the person referred to in *p*. We may ask then, given the situation Sam is in, as an epistemic agent what *ought* she do? The first sort of activity over which we have control, and which is closely tied to responsible agency, is whether *to engage* in the exercise of deliberations or not. This is not to say that we are at all times obligated to deliberate whenever a possible belief arises. Most of our non-inferential beliefs, like perceptual or memorial beliefs, are justified, yet we very infrequently need to actively scrutinize our perceptual or memorial evidence every time we come to a belief. In other kinds of cases, due the complexity of the subject matter or of one's own evidential situation, the agent is handed the obligation to engage in deliberations. If Sam's is one of these situations, and she bypasses the procedure altogether and just goes ahead and believes *p*, she fails to meet this epistemic obligation. We have an expression in English that seems to capture this kind of failure, as when we accusingly reproach others by saying, "You just weren't thinking hard enough."

Suppose Sam meets her obligation here and chooses to engage in deliberations. Insofar as deliberations are activities used to come to an answer to

the question of whether  $p$ , there is very little control we have as to where the deliberations will end up.<sup>23</sup> Deliberating whether or not Sam comes to believe  $p$  is entirely a matter of what it is that she finds compelling. But, surely, neither she nor anyone in like situations have any choice about what she finds compelling; what is compelling to one is entirely a matter toward which one's reasons and non-epistemic factors point.

Sometimes the mere act of beginning deliberations is not enough to exclude irresponsible epistemic conduct. In the action case, one may know the general rule that fighting is bad, but when deliberating what to do on a specific occasion, it may be that one's temptation to fight the enemy is too powerful to keep oneself from fighting. In a similar way, while most of us know that it is a good thing in general to base beliefs exclusively on our evidence, our deliberations are sometimes overpowered by some strong passion or temptation (e.g. a wish that  $p$  be true) that tend to steer beliefs away from the truth. If M1 is one of those strong temptations for Sam, it may be that her deliberations are already at the very start geared toward where that temptation points. There is nothing Sam can do about this fact. What, then, is it that Sam and other similar such agents are supposed to do? I propose that they *ought to do those mental actions that tend to facilitate a proper response to the acquired evidence*. Notice the difference between the 'ought' claim listed here and Feldman's OUG. I am not saying that our obligations are to properly respond to the evidence. The reason I say this is that what we respond to in deliberations, what mental states upon which we wind up basing our beliefs, is something over which we have very little control. I cannot just decide to maintain one of my currently held beliefs but change the basis for which I hold it. I cannot directly alter the connections that bind my beliefs and other mental states within the space of reasons. Demanding that one respond properly to one's evidence is analogous to telling Timmy to be in a sleeping state at this very moment, which as I have argued is something he cannot do.

Just as there are bodily actions Timmy can do so as to facilitate the state of sleep (e.g. lie down in bed), there are mental actions epistemic agents like Sam can do so as to facilitate properly responding to one's evidence. In particular, while in deliberations one can willfully choose to pause and ask (or even re-ask) oneself the question: "What are my reasons here?" The very asking of this question prompts the agent to change course and take a higher-order, reflective stance towards one's own doxastic situation and, moreover, determine what factors have been influencing deliberations up until the question was asked. In taking this stance,

---

<sup>23</sup> This is a point contested by Steup ("Doxastic Voluntarism and Epistemic Deontology" and "Doxastic Freedom").



the agent is more likely to be caused to redirect her mind away from her non-epistemic temptations or passions and to zero in on the evidence. “Are these reasons any good?” is another sort of self-reflective question that, in the very activity of asking it, brings the agent to scrutinize her evidence, ensuring a higher likelihood that whichever belief she ultimately forms is the proper response from her evidence.

As I say the sorts of mental activities mentioned here are those that tend to lead to proper evidential responses. They do not always do so. If M1 is Sam’s wish that *p* be true, the temptation may be so powerful that even after she has asked herself the self-reflective questions, her belief that *p* is ultimately a response to that wish (just as the mother who cannot accept the lead detective’s testimony that her son is dead due to her psychologically compulsory hope that he is still alive). From a deontological point of view, I claim that Sam has nonetheless fulfilled all of her epistemic obligations. She is a responsible agent.

There is an objection to the account I am proposing, one which comes from Alston in his seminal paper.<sup>24</sup> While he concedes that we do possess a kind of indirect voluntary influence over our beliefs, of the sort that I have been describing here, Alston ultimately concludes that deontological concepts like requirement and blame, as tied to this indirect sort of influence, cannot underwrite the concept of epistemic justification. For, he thinks that this deontological approach to epistemology fails to capture the truth-conducive nature of epistemic justification: epistemically justified beliefs are true or at least likely to be true. The example with Sam above illustrates this complaint well. Sam, as I have claimed, stands in no violation of any epistemic obligations, and yet beliefs which are formed in such a way on the basis of a wish tend to be highly unreliable. So, assuming that Alston is correct about the objective relation between epistemic justification and truth, are we forced to abandon epistemic deontology altogether?

I argue that the answer is “no.” For it is consistent, on the view that I am advancing, that an agent can act (epistemically) responsibly with respect to the belief her actions influence while at the same time falling short of acquiring a (epistemically) justified belief. With regard to Sam, then, it may well be that she is unjustified in believing *p*, given that it was ultimately grounded in a wish. But there are no grounds for blaming her for how she performed. (Similarly, even if Timmy winds up having a completely sleepless night, his parents cannot rightly blame him for this, given that he did all the necessary actions that tend to bring

---

<sup>24</sup> Alston, “The Deontological Conception,” Section 7.

about sleep.) How, then, can it be that an agent could be responsible for holding a belief which she is nevertheless unjustified in holding?

I will close with one brief suggestion, and to do so let me first introduce a distinction from the literature that often gets made in the context of responsible epistemic agency.<sup>25</sup> We may make, on the one hand, an *evaluative* epistemic assessment of the *belief* formed by the agent, according to which beliefs are assessed favorably as far as they satisfy some epistemically relevant good. Thus, along reliabilist lines, for example, a belief receives a favorable evaluative assessment insofar as the belief was produced by what is in point of fact a reliable process. Or, along evidentialist lines, we would say that a belief's evaluative assessment is determined by how well the belief fits with the agent's evidence at the time. On the other hand, we may instead offer a *regulative* epistemic assessment of the *agent* with respect to the belief she formed. This sort of assessment is determined strictly in accordance with the agent's own point of view. Taking stock of everything present to the agent's mind at the time, including not only the justification-conferring grounds or evidence of which she is in possession, but also all doxastically relevant factors – like emotional or other non-epistemic influences and her ability to resist them in forming a judgment, her competency in reasoning, and her ability to recognize rational links between her reasons and the proposition her reasons support – the primary question is a deontological one: has she done all that can be expected of her to do?

With respect to the question, “in virtue of what are agent's beliefs epistemically justified?” a strong internalist answer would be that justification is a matter of receiving favorable regulative epistemic assessment. For a strong externalist, in contrast, justification is a matter of receiving a favorable evaluative epistemic assessment. But neither of these assessments, on their own, are able to capture that special feature of justification we all deem to be so important, namely, that a justified belief is such that a) it bears an objective likelihood of truth and b) the agent who forms this belief has done so in a responsible manner. Therefore, I suggest that the concept of justification is sufficiently complex, such that deontological assessments are but one component, out of multiple components, to

---

<sup>25</sup> The sort of distinction I have in mind here has been made in places such as Alvin Goldman, “The Internalist Conception of Justification,” *Journal of Philosophy* 75 (1980), Kent Bach, “A Rationale for Reliabilism,” *The Monist* 68 (1985), Philip Goggans, “Epistemic Obligations and Doxastic Voluntarism,” *Analysis* 51 (1991), and Alvin Plantinga, *Warrant: The Current Debate* (Oxford: Oxford University Press, 1993), 15-19. See also James Pryor, “Highlights of Recent Epistemology,” *The British Journal for the Philosophy of Science* 52 (2001): Section 4 for a discussion.

determining whether an agent is justified.<sup>26</sup> I myself hold that the demands of justification are such that being justified implies that the agent's doxastic performance receives favorable assessments from *both* a regulative and an evaluative perspective. That is, if S's belief that p is justified, then it is true, first, that she has fulfilled all her deontological requirements, and secondly, that the belief formed satisfies some epistemically relevant good, such as its actually and objectively being a proper response to the agent's evidence. Thus, while it may be true that all justified agents are responsible agents, it does not follow that all responsible agents are justified agents.

## 5. Conclusion

In conclusion, I have argued that Richard Feldman's 'role-ought' response to Alston's argument from doxastic involuntarism fails. However, when we find those roles that are most analogous to that of being a believer, we find that there is room to be made for deontological assessments even if we lack any suitable form of doxastic control and also that "ought" does imply "can." The upshot is that the deontological dimension of justification requires, not voluntarily forming any particular belief, but voluntarily performing those mental actions that tend to bring about a self-reflective stance in the agent, which in turn increases the likelihood that the agent's doxastic state properly connects back to her evidence.<sup>27</sup>

---

<sup>26</sup> See fn 3.

<sup>27</sup> For their helpful comments on earlier versions of this paper, I wish to thank Kevin Falvey, Anthony Brueckner, Michael Rescorla, Joe Berendzen, Ben Bayer, Patrick Leland, the Epistemology Research Group at the University of Edinburgh, and participants at the first annual Southern Epistemology Conference.