

EPISTEMIC DUTIES AND FAILURE TO UNDERSTAND ONE'S EVIDENCE

SCOTT STAPLEFORD

St. Thomas University

Abstract. The paper defends the thesis that our epistemic duty is the duty to proportion our beliefs to the evidence we possess. An inclusive view of evidence possessed is put forward on the grounds that it makes sense of our intuitions about when it is right to say that a person ought to believe some proposition *P*. A second thesis is that we have no epistemic duty to adopt any particular doxastic attitudes. The apparent tension between the two theses is resolved by applying the concept of duty to belief indirectly.

Keywords: Epistemic duties; ethics of belief; doxastic voluntarism; evidence; evidentialism; R. Feldman.

Richard Feldman argues that we never have an epistemic duty to believe any proposition *P* that our evidence supports if we fail to apprehend the connection between our evidence and *P*.¹ I argue that an implication of this view is that we never have an epistemic duty to believe any proposition that we do not in fact believe. My argument depends upon the falsity of one fairly implausible version of doxastic voluntarism. Given that this conclusion conflicts with our understanding of what it is to have a duty, I infer that the only epistemic duty we have is the duty to proportion our beliefs to the evidence, a view which is consistent with our never having an epistemic duty to believe some proposition *P*.

The idea that I want to defend is suggested by a passage in John Locke's *Essay Concerning Human Understanding*. In Chapter 15 of Book IV, Locke appears to make the obtaining of due proportion between one's beliefs and one's evidence a requirement of rationality:

[T]he mind if it will proceed rationally, ought to examine all the grounds of probability, and see how they make more or less, for or against any probable proposition, before it assents to or dissents from it, and upon a due balancing the whole, reject, or receive it, with a more or less firm assent, proportionably to the preponderancy of the greater grounds of probability on one side or the other (Locke 1975: IV.xv.5: 656).

The thought here can be formulated as an epistemic principle governing belief. We'll call it Locke's 'principle of proportionality' or 'PBE':

Principia 16(1): 147–177 (2012).

Published by NEL — Epistemology and Logic Research Group, Federal University of Santa Catarina (UFSC), Brazil.

PBE: We have an epistemic duty to proportion our beliefs to the evidence we possess.²

It is not a stretch to think that acceptance of the proportionality principle would necessitate the recognition of a second sort of epistemic duty, the duty to adopt particular doxastic attitudes, where having a duty to adopt some doxastic attitude is understood as follows:

S has an epistemic duty to adopt some doxastic attitude *D* at t_1 iff there is some proposition *P* such that it is *S*'s epistemic duty to believe, disbelieve or suspend judgement about *P* at t_1 .

The line of reasoning leading from acceptance of PBE to acceptance of a second type of duty would go something like this: An agent whose doxastic attitudes fit her evidence perfectly would be in a cognitive situation such that her duty to obtain proportion between her beliefs and her evidence would be fully satisfied. And adopting all and only fitting doxastic attitudes would of necessity bring about a state of perfect proportion between belief and evidence. Therefore, there are particular doxastic attitudes—the ones that fit our evidence—that it is our epistemic duty to adopt. On this account, a corollary of PBE is the principle ADA:

ADA: We have an epistemic duty to adopt certain doxastic attitudes.

Though my endorsement of Locke's proportionality principle appears to commit me to the idea that we have duties to adopt certain doxastic attitudes (those that fit our evidence), there is no need, I will argue, to suppose that this is so. The two main theses of this paper are that PBE is true and that ADA is false: We never have a duty to believe, disbelieve or suspend judgement about any proposition *P*, though we do have an epistemic duty to proportion our beliefs to our evidence.³

My first goal is to show that ADA is false. I attempt this in Section 1. In Section 2, I draw a distinction between two types of evidence and argue, first, that we have as evidence more than is sometimes supposed and, secondly, that we should deflate 'oughts' as they apply to particular propositions. This will allow me to argue in Section 3 that a certain version of PBE is true, though my understanding of this principle is probably quite diverse from Locke's.

1. There are No Duties to Believe

Arguments against the existence of epistemic duties sometimes appeal to the falsity of doxastic voluntarism, a thesis of empirical psychology according to which belief is within our voluntary control. William Alston, who favours this approach, examines a number of versions of doxastic voluntarism and argues that all of the epistemically

significant forms are false. Actions that we can perform through a simple act of will are within what he calls our *basic voluntary control* (Alston 1988: 263–68). I can raise my arm, for instance, by simply willing it, but I cannot just will my teacup to rise. Actions that we can perform straight away by doing something else are within our *non-basic immediate voluntary control* (Alston 1988: 268–74). I can lift my teacup right now by raising my arm, but I cannot do anything on the spot to cause a wardrobe to exist. There are also things that we can do voluntarily over time by doing other things—they fall within our *long range voluntary control* (Alston 1988: 274–77). I could eventually build a wardrobe by working with wood, but I cannot take any course of action that would result in the moon reversing its orbit. Alston considers a fourth sort of voluntary control as well, but we can safely ignore it here. The relevant point is that we do not appear to exercise voluntary control over our beliefs in any sense that matters epistemically. I cannot believe that there is a goat under my desk right now by simply willing it (this would be an example of *basic doxastic voluntary control*),⁴ nor can I make myself believe that my colleague (whom I see across the hall) is not in his office by simply shutting my door (an example of *non-basic immediate doxastic voluntary control*). Likewise, there is no long range course of action I can take that would cause me to believe that my colleague can levitate (an example of *long range doxastic voluntary control*).

On the assumption that these forms of doxastic voluntarism are false, it is sometimes asserted that the deontological conception of epistemic justification—according to which there are things that we epistemically should or shouldn't do—is confused and that we therefore have no epistemic duties. In outline, the argument is this:

The Standard Voluntarism Argument

- (1) Deontological judgements about beliefs are true only if we have voluntary control over our beliefs.
- (2) We do not have voluntary control over our beliefs.
- (3) Deontological judgements about beliefs are not true.⁵

One possible response to this argument is to offer grounds—if there are any—for thinking that belief is voluntary after all and that premise (2) is therefore mistaken.⁶ In Section 3, I will explicate our basic epistemic duty in terms that rely on our capacity to control some of our beliefs now and over time by doing other things (by exercising a very specific sort of *non-basic immediate* and *long range voluntary control*), but I do not think that this is needed in order to defeat the standard voluntarism argument. Even if we had no control over belief whatsoever, this would do nothing to show that belief is beyond epistemic appraisal. Feldman argues plausibly that premise (1) is false, that deontological judgements apply correctly to beliefs whether

or not beliefs are subject to voluntary control. He develops the useful idea of role oughts. Role oughts are requirements making reference to standards of good performance that apply to a person insofar as the person plays a certain role. In Feldman's view, there are certain things that we ought to do just insofar as we occupy the role of believers: "[W]e form beliefs in response to our experiences in the world. Anyone engaged in this activity ought to do it right . . . It doesn't matter that in some cases we are unable to do so" (Feldman 2004a: 175). Conee and Feldman also contrast evaluation of belief with evaluation of the person. If a person spontaneously forms the belief that the lights are on in a lighted room, their belief is justified, and hence subject to epistemic evaluation, whether or not the person has any control over the belief. It might be unfair to evaluate the *person* for forming beliefs correctly or incorrectly if doxastic voluntarism is false, but it is certainly not unfair to evaluate *beliefs* (Conee and Feldman 2004: 84–6).

These thoughts cast serious doubt on the cogency of the standard voluntarism argument against epistemic duties. The falsity of some or all forms of doxastic voluntarism in no way implies that beliefs are immune to epistemic evaluation. On the contrary, beliefs are subject to criticism just insofar as there are things that we ought or ought not to believe on the basis of our evidence. But the falsity of *basic* doxastic voluntarism—which I assert as evident—backs a key premise in my argument that we have no duty to adopt any particular doxastic attitudes.

The argument rests on three crucial assumptions for which I can provide only modest support here (though more support along the same lines can probably be found). The first assumption is this: If *P* is logically supported by *S*'s evidence and *S* fails to grasp the connection between her evidence and *P*, then *S* has no *duty* to believe *P*.⁷ The idea comes from Feldman, who says that in cases where a person has evidence implying some proposition such that the connection between their evidence and the implied proposition is distant and difficult to see, the person ought not believe the proposition (it would be "rash" to do so, barring other reasons for belief) (Feldman 2004a: 181). Though I think that there is some sense of 'ought'—albeit a weaker sense—in which the person *ought* to believe even this remote logical consequence of her evidence, it seems to me that it would be too demanding to say that she has a *duty* to believe it.⁸

Suppose that *S* believes *Q*, *R* and *T*, and that a long sequence of logical deductions, far beyond *S*'s immediate comprehension, leads from the conjunction of *Q*, *R* and *T* to *P*. Surely we don't want to say that *S* has a duty to believe *P* if she is currently blind to the logical connection that her evidence bears to *P*. I do not know how to support this claim any further than by appealing to favourable intuitions and by stressing that I also hope to accommodate the countervailing intuition that the logical connections themselves have some claim to be binding on *S* despite her failure to grasp them. The bindingness of logic will be prominent in Section 2. Here

I just want to emphasise that the language of 'duty' is arguably too strong to apply directly to the beliefs of an agent who fails to understand what her evidence implies. This is the first premise of my argument.⁹

The next assumption I make is tied to the falsity of basic doxastic voluntarism. It seems evident from introspection that where a person does grasp the connection between her evidence and some proposition supported by that evidence,¹⁰ she will spontaneously believe that proposition (I assume that grasping an evidential connection involves considering the proposition that is evidentially supported). Cases where apprehension of an evidential connection is clouded by emotion or blinding conviction should be counted as cases of failure to fully understand one's evidence. So if I have a clear apprehension of the connection between my evidence and some proposition *P* that is supported by that evidence, I will believe *P*.¹¹ To take an example, suppose that my evidence overwhelmingly supports the proposition that my teapot is not at this moment empty: I filled the pot one minute ago, poured one cup, checked twice to see that the pot was still more than half full, and never let it out of my sight since I replaced the lid.¹² Suppose I also see clearly that my evidence supports the proposition 'The teapot is not empty' (again, this involves actually considering the proposition). In this situation, I think it is a matter of psychological necessity that I will believe, 'The teapot is not empty'. I just don't command basic voluntary control over my beliefs.

In general, a person cannot fail to believe *P* when *P* is supported by their evidence and they see clearly that it is.¹³ I want to stress once more that a person who for emotional reasons 'just can't' believe some proposition supported by their evidence is not, in my view, actually comprehending the evidence. If they could only be made to perceive the strength of the reasons they have for accepting the proposition in question, assent would follow at once. We don't always understand our evidence, but when we do, belief tracks it spontaneously.¹⁴

The final assumption is that duty is the sort of thing a person can fail to meet. This follows directly from the concept of duty. On the standard voluntarism argument, the worry was that we perform *or* fail to perform our alleged duties involuntarily, and epistemic evaluation was thus seen to be somehow unfair or out of place. The idea here is somewhat different. It just seems to be part of the notion of duty that whenever something is required it is not something that happens without exception. You have no pneumatic or gastronomic duties to breathe or to digest, since these are not actions that you can ever omit. Assuming well-functioning motor nervous, muscular and skeletal systems, you have no kinetic duty to move your limbs, since you cannot fail to move your limbs at least sometimes. You have no soporific duty to fall asleep at intervals, since sleeping occurs in cycles by a natural necessity (it can be postponed, but not indefinitely). By contrast, you do have a legal duty to register your car, and this is something that you can fail to do. You have a moral duty

(I assume) to help those in need, but you might not do it. You have a contractual duty to pay your mortgage, though it may be beyond your power to do so.¹⁵ It seems to be universal that any action we classify as a duty will be an action that we can fail to perform, and any action that we perform without omission will be an action that we do not regard as a duty. These examples provide at least some support for the view that duties are the sorts of things that occasionally go unfulfilled.

Some further clarification of this final assumption may be helpful. The point is not that belief must be within our voluntary control for it to be open to epistemic evaluation. Perhaps none of our actions are ever within our effective voluntary control. The point is rather that only where two contrary outcomes are possible does the concept of duty apply. It may be the case that if I cannot fail to believe *P*, then believing *P* is not within my voluntary control, since voluntary control “necessarily extends to contraries,” as Alston says (1988: 261). But even if that were true, the lack of basic voluntary control is not what makes duty the wrong notion here. At my university, we have an institutional duty to submit course outlines to the Registrar’s Office before classes begin. If I fail to do this—it happened once—then I have violated one of my institutional obligations. But if all actions are involuntary (in the sense of being fully determined by causal antecedents that are themselves beyond our voluntary control), then both submitting and not submitting outlines are beyond my voluntary control. Yet I still have a duty to submit them. By contrast, I have no institutional duty to lecture while conscious, since I cannot fail to be conscious while lecturing.

One could still urge that this defence of my third assumption relies on a version of the standard voluntarism argument, since the difference between, say, submitting course outlines (a required action) and breathing (an exceptionless non-required action) is that the former is something we *will* while the latter is not, even if neither action is ‘free’ or uncaused. On this understanding of my argument, the deontological conception of duty would not apply to belief because belief is not something that we will, just as in the standard voluntarism argument the concept of duty does not apply because belief is not within our voluntary control. I have three responses here. First, my ‘kinetic’ example above makes clear that there are non-required actions (moving our limbs) that we perform voluntarily (in the sense of willing them) *and* without fail, so it can’t be that the essential difference between required and exceptionless non-required actions is that the latter are all involuntary in the sense of being non-willed. Secondly, my argument would not go through without the first assumption, for if there could be a duty to believe some proposition *P* where we fail to see the connection between our evidence and *P*, then we could fail to perform our duty by failing to understand the relevant connection and failing to believe *P* as a result. The standard voluntarism argument makes no such assumption. Thirdly, even if I concede (there’s no reason not to concede) that my refutation of ADA is a

special type of voluntarism argument since it presupposes the falsity of *basic* doxastic voluntarism, it differs from the standard voluntarism argument in three key respects: (1) I argue in Section 2 that deontological judgements involving epistemic oughts are true whether or not belief is voluntary in any sense (contrary to the first premise of the standard voluntarism argument); (2) I claim in Section 3 that belief is subject to non-basic voluntary control in an epistemically relevant sense (contrary to the second premise); (3) I also contend in Section 3 that we do have an epistemic duty with respect to belief after all (contrary to the conclusion).

Returning now to the main argument, it should be obvious where all of this is heading. If we fail to see the connection between our evidence and some proposition *P* supported by that evidence, then we have no epistemic duty to believe *P*. An implication of this is that if we do see the connection between our evidence and some proposition *P* which is evidentially supported by it, there is no duty to believe *P* in that case either. For if we do see the connection between our evidence and *P*, then we will believe *P* (assuming that basic doxastic voluntarism is false). And we have a duty to believe a proposition only if we see the connection between the proposition and our evidence (since if we do not see the connection, then we do not have a duty to believe it). So we have a duty to believe some proposition *P* only if we do in fact believe *P*. But then it makes no sense to speak of duty here, since we have a duty to do something only if we could fail to do it. Therefore, we have no epistemic duty to believe any particular proposition.

The same reasoning applies to the doxastic attitudes of disbelief and suspension of judgement. Using the notion of a fitting doxastic attitude, where the fitting attitude for a person to adopt is the one that is supported by the evidence the person has, the argument can be stated more precisely as follows. (I make the common assumption that belief is a type of action, though nothing essential hinges on this.)

The Argument against ADA

For any subject *S*, any proposition *P*, any doxastic attitude *D*, any action *A* and any time *t*:

- (1) If *S* fails to understand that her evidence at *t* makes *D* the fitting doxastic attitude for *S* to have toward *P* at *t*, then it is not the case that *S* has an epistemic duty to have *D* toward *P* at *t*.
- (2) If *S* understands that her evidence at *t* makes *D* the fitting doxastic attitude for *S* to have toward *P* at *t*, then *S* cannot fail to have *D* toward *P* at *t*.
- (3) *S* has an epistemic duty to have *D* toward *P* at *t* only if *S* understands that her evidence at *t* makes *D* the fitting doxastic attitude for *S* to have toward *P* at *t*. (1) [by contraposition]

- (4) *S* has an epistemic duty to have *D* toward *P* at *t* only if *S* cannot fail to have *D* toward *P* at *t*. (2), (3) [by hypothetical syllogism]
- (5) *S* has a duty to perform *A* only if *S* can fail to perform *A*.
- (6) If *S* cannot fail to perform *A*, then *S* does not have a duty to perform *A*. (5) [by contraposition]
- (7) *S* has an epistemic duty to have *D* toward *P* at *t* only if *S* does not have an epistemic duty to have *D* toward *P* at *t*. (4), (6) [by hypothetical syllogism and replacement]
- (8) *S* does not have an epistemic duty to have *D* toward *P* at *t*. (7)

Premise (1) is the first assumption that we discussed above. Premise (2) is the second assumption, and premise (5) is the third. Premise (7) says that having a duty to adopt a particular doxastic attitude implies *not* having a duty to adopt a particular doxastic attitude. This can only be the case if we do not have a duty to adopt a particular doxastic attitude. And so we infer the falsity of the antecedent and draw the conclusion (8), which is general and equivalent to the negation of ADA.

If ADA is false, as our argument has shown, there may be some difficulty in explaining the seeming correctness of certain judgements we routinely make about other people's epistemic mistakes. We say that *S ought* to believe *P*, given her evidence, when *P* is a logical consequence of other things that *S* holds to be true, even when *S* is unaware of the logical relations that hold between *P* and her evidence. And we say that *S ought not* believe *Q*, if *Q* is logically inconsistent with *S*'s other beliefs, whether or not *S* perceives this. What force can such judgements have, if there are no epistemic duties to adopt any particular doxastic attitudes? At the end of the next section I make a preliminary case for *deflating* epistemic oughts as they apply to particular propositions: Judgements such as '*S ought* (or *ought not*) to believe *P* on the basis of her evidence' are legitimate, but their force is only to express a certain sort of normativity without implying any duties to believe particular propositions. How widely deflated epistemic oughts apply depends essentially on what we include in the evidence a person has available to them at a given time. The question of the evidential availability of the rules of logic is of central importance here, and it is our focus in the discussion to follow.

2. Two Types of Evidence

Feldman thinks that at any given time we have as evidence only those propositions that we are actually thinking of at that time (Feldman 2004b: 219, 232ff.) I will argue, contrary to Feldman, that in order to make sense of the notion of an epistemic ought in a way that squares with our intuitions about when it is right to say

that someone is believing something irrationally, we need to draw a distinction between the evidence that a person has cognitively accessed—the propositions the person is actually thinking of—and the evidence that is cognitively accessible to that person—the propositions that the person could access without engaging in any additional evidence-gathering activities. The total evidence that a person has available to them at a given time comprises their cognitively accessible evidence and their cognitively accessed evidence taken together—‘CAB evidence’ and ‘CAD evidence’, respectively.¹⁶

In many cases it will be easy to say whether or not some factual proposition *P* belongs to one's evidence and, where it does, whether it belongs to one's CAB evidence or to one's CAD evidence. Since I know the proposition ‘The University of Marburg was founded in 1527’ to be true, it belongs to my CAD evidence whenever I think about it and to my CAB evidence whenever I do not. The proposition expressing Martin Luther's date of birth, on the other hand, does not belong to my evidence at all. While I know several pertinent facts about Luther's life, I do not know the exact date on which he was born. So I do not know the truth-value of the following proposition (it's a guess): ‘Martin Luther was born on March 18th 1482’. No matter how much I think about it, I will not be able to ascertain its truth-value—and so add either it or its negation to the set of propositions that are available to me as evidence—without acquiring some additional information about the world through sensation (whether by reading, watching a documentary, asking a colleague, or engaging in some other evidence-gathering activity). Because new sensory input is needed in order for me to determine the truth-value of ‘Martin Luther was born on March 18th 1482’, the proposition does not belong to my CAB evidence. Most factual propositions are like this: If we have learned them, and can recall them, they belong to our CAB evidence. If we have not learned them, or cannot recall them, they do not.

As a first and rough approximation we could therefore delimit our notions as follows:

CAD *P* belongs to *S*'s CAD evidence at t_1 iff *S* is thinking of *P* at t_1 .

CAB *P* belongs to *S*'s CAB evidence at t_1 iff *P* belongs to *S*'s CAD evidence at t_1 or *S* could access a memory of *P* at t_1 without relying on any additional sensory input.¹⁷

One obvious problem with this account of CAB evidence is that there are varying degrees of accessibility for beliefs stored in memory: Some memories are recent and clear; others are distant and obscure. How distant and obscure a memory needs to be before it falls outside the scope of what is cognitively accessible to a subject is far from clear.¹⁸ But I want to set this problem aside for now and consider the problem of how to classify a particular subset of non-factual propositions.¹⁹ The above characterization of CAB evidence does not say anything about non-factual

propositions such as the truths of logic and mathematics. It does seem possible, at the very least, that certain non-factual propositions are available as evidence to a subject *S* even though *S* has never acquired knowledge of them through experience nor even thought about them. I want to focus here on the truths of logic,²⁰ though what I say may well apply to the basic truths of mathematics as well.

It is not easy to classify propositions expressing truths of logic and the relations of logical support that obtain between propositions.²¹ Does one need to be taught the rules of logic in order for them to be available as evidence? I will argue that the basic rules of logic and good thinking in general are cognitively accessible to every rational agent and that they are thus a part of every rational agent's CAB evidence. To see why this is so, it will be useful to start with an example.

Suppose some first-year philosophy student, Holly, has learned and accepted the rules of inference called *Modus Ponens* and *Modus Tollens* and is not currently thinking about them. The propositions expressing these rules, therefore, belong to Holly's CAB evidence. But what about the rule of inference, *Hypothetical Syllogism*, which was not covered in the course? Surely this rule is also available to Holly just insofar as she is rational, even though it was not brought to her attention by a philosophy instructor. It wasn't the mere fact that she was taught *Modus Ponens* and *Modus Tollens* that made her accept them, after all, but the fact that she could just see (in some intellectual sense of 'see') that they were good rules. The rules of logic—including the ones that she has never been taught—are a part of Holly's cognitively accessible evidence because she can understand that they hold just by thinking about them. I think I can offer an argument to show that this is in fact the case, that we need a distinction between CAB evidence and CAD evidence, and that the rules of logic—at least the basic ones falling under the heading 'rules of good thinking in general'—are cognitively accessible to every rational agent (that is to say, that they belong to everyone's CAB evidence).

The argument appeals to our intuitions about when it is right to say that adding a certain belief to a subject's belief set would cause an inconsistency within the set. Consider two subjects, S1 and S2. Suppose that S1's belief set includes the belief that $P \ \& \ Q$, where 'P' and 'Q' stand for any two factual propositions. And suppose that S1 considers adding the proposition ' $\sim Q$ ' to her belief set. We would say that the proposition ' $\sim Q$ ' is contrary to her evidence and that adding it to her belief set would create an inconsistency. And I think that we could say this without possessing any knowledge of the other propositions that S1 believes. We must be assuming, therefore, that the conjunction elimination rule and the principle of non-contradiction are a part of S1's cognitively accessible evidence. If they were not a part of her evidence, then adding ' $\sim Q$ ' would cause no inconsistency in her belief set. ' $\sim Q$ ' would not conflict with her evidence, since the conjunction elimination rule would not be available to her, and so no contradiction could be generated between

' $\sim Q$ ' and her belief that $P \& Q$. We *do* say that a contradiction would result from adding the belief ' $\sim Q$ ' because we *do* think that the laws of logic are accessible to S1 and are therefore a part of her CAB evidence.

Now consider S2, whose belief set, we'll suppose, contains the proposition 'Angela Merkel lives in Hamburg.' Imagine that S2 is considering the proposition 'The current chancellor of Germany lives in Berlin.' Would adding this proposition to S2's belief set be contrary to her evidence? A moment's reflection reveals that we cannot answer this question until we acquire some information about the other factual propositions that S2 holds to be true. Specifically, we need to know whether S2 believes that Angela Merkel and the current chancellor of Germany are the same person. We don't automatically attribute a belief concerning the identity of Merkel and the Chancellor to S2—since it is a factual matter—and so we don't automatically assume that the proposition asserting the Chancellor's address conflicts with S2's evidence. But in the case of S1 we would say that adding the proposition ' $\sim Q$ ' to her belief set would conflict with her evidence *without* knowing her other factual beliefs.

My point is that we see a conflict in the first case between the proposition considered and the subject's evidence because we assume that the laws of logic are accessible to her and that they are thus a part of her CAB evidence. We do not make any such assumptions about factual propositions. We do not say that adding the belief that the Chancellor lives in Berlin conflicts with S2's evidence unless we know that S2 believes that the Chancellor and Angela Merkel are the same person. We do not just assume knowledge of this factual identity on the part of S2 because knowledge of this or that factual proposition is just not the sort of thing that we can reasonably attribute to rational agents *a priori*. But our intuition that there is a conflict between the proposition ' $\sim Q$ ' and S1's evidence (which includes the belief that $P \& Q$) suggests that we do suppose that the laws of logic are a part of every rational agent's cognitively accessible evidence. I suggest that our intuition in this case is right. Now I want to try to strengthen that intuition with an argument that makes use of an example which Feldman cites in support of the opposite conclusion.²²

Feldman considers a report by Tversky and Kahneman on a series of psychological experiments that seem to reveal systematic irrationality in test subjects. In one experiment, the subjects were given a description of a typically liberal and politically active woman and asked which of two categories the woman is more likely to be in. Many subjects answered, incorrectly, that it is more likely that the woman is a feminist bank teller than a bank teller, evidently focusing on representative characteristics of liberal and politically active women and ignoring the basic rules of probability theory (Feldman 2004b: 221). Given the conjunction rule of probabilities—which says that the probability of a conjunction (feminist and bank teller, here) cannot exceed the probability of one of its conjuncts (bank teller)—the correct inference to make

from the description provided is that it is more probable that the woman is a bank teller. The conjunction rule is simple and intuitive, but test subjects who give the wrong answer are most likely misled by the fact that being a feminist is typical of a liberal, politically active woman. In Feldman's view, if a subject is either not aware of the conjunction rule or not thinking of it, then it might actually be rational for them to conclude that it is more likely that the woman is a feminist bank teller than a bank teller based on the association they are making between active liberalism and feminism (Feldman 2004b: 222). This follows from the fact that Feldman limits available evidence to the propositions one is actually thinking of.

This implication of Feldman's position is in my view an argument against it. If we deny that the rules of logic are a part of everyone's CAB evidence, it seems to follow that there will be cases like the one described above, where it is rational for a person to believe a proposition P though P is not logically supported by the person's evidence. This is tantamount to denying the normativity of logic. My point can be put this way: If we acknowledge that logic is normative—in the sense that it tells us how we *ought* to think—then Feldman's position will require us to accept that sentences (1) and (2) below both accurately characterise the epistemic situation of the test subjects who answered incorrectly (where 'P' stands for 'It is more probable that the woman described is a feminist bank teller than a bank teller'):

- (1) Test subject S *ought not* infer P at t_1 on the basis of her evidence.
- (2) It is *rational* for test subject S to infer P at t_1 on the basis of her evidence.

It seems to me that sentences (1) and (2) are inconsistent and that Feldman's theory commits us to both of them, so long as we accept that P is the *wrong* answer according to logic. Objectively speaking, it is *not* more probable that a liberal and politically active woman is a feminist bank teller than a bank teller (since it is not more probable that $A \& B$ than it is that A , whatever 'A' and 'B' stand for). Sentence (1) is true—the subject *ought not* infer P from the description—because P is the wrong answer.²³ Logic tells us that from the evidence provided to the test subjects, one *should not* draw the conclusion P . Given that P is not warranted by the description, the test subjects *epistemically ought not* infer P . (This is precisely why the study has been taken by some to reveal systematic irrationality.) But it will be simultaneously *rational* for the test subjects to infer P —sentence (2) will be true—if we deny that the rules of logic are available to them as evidence. On the contrary, I do not believe that sentences (1) and (2) can be true simultaneously. My belief that logic is normative and that the test subjects should therefore *not* infer P commits me to the truth of (1). Since I see no good reason to deny that the rules of good thinking are binding on the test subjects—and since exemption by ignorance is what supposedly makes it rational for them to draw the wrong conclusion—I infer that (2) is false. It is *not*

rational for the test subjects to conclude that the woman is more likely a feminist bank teller than a bank teller on the basis of the description provided to them.²⁴

True, the test subjects' CAD evidence that the woman is politically active and liberal, combined with their failure to see the relevance of the conjunction rule, certainly makes it *understandable* that they infer P from the description—the frequency with which feminism occurs in connection with liberalism and political activism is a potent red herring—but I think we must deny that it makes it *rational* for them to do so. Given the normativity of logic—given that, logically speaking, P is the wrong answer—we must say that they should not infer P from the evidence. This just means that the rules of logic are binding on them and that they are thus in some sense cognitively accessible to them. (One can't evade the 'oughts' of logic by neglecting them or applying them incorrectly.)

Keeping our focus on logic, I think we can go further and say that not even the holding of *false* beliefs about logic can justify one in making illogical inferences from evidence possessed. The situation is very different with factual propositions, of course, where false beliefs can lead justifiably to other false beliefs. A false belief about logic, however, can never make it rational for a subject to draw a conclusion for which there is no logical warrant. Examples should make this clear. Suppose that Holly encounters a horse for the first time, and the horse happens to be white. And suppose that Holly has an idiosyncratic belief about enumerative induction according to which one instance is a good sample size from which to make a universal generalization. As a result, Holly infers that all horses are white. Though Holly's faulty logic seems to allow this inference, her argument is inductively weak nonetheless. If we agree that logic codifies the rules of good thinking and that only good thinking is rational, then we must agree that Holly's inference is irrational. Similarly, if Holly believes

(1) $P \supset Q$

and

(2) $\sim P$

she is not justified in inferring from this evidence

(3) $\sim Q$

even if she also believes that the (bad) rule of inference 'denying the antecedent' is logically valid. Finally, to take a less contrived example, suppose that Holly has read many articles by reputable scientists in mainstream newspapers and magazines stating that man-made global warming is a major threat to the environment and presenting some of the most compelling evidence for its reality. Suppose that Holly

then reads a leaked email exchange between two climatologists in which ‘tricks’ for fudging data in favour of global warming are proposed as a means of silencing sceptics. Whatever Holly’s beliefs about logic may be, it is not rational for her to infer from this correspondence alone that the threat of global warming is not real, that vast numbers of conspirators are colluding to skew the data and that the entire scientific enterprise is hopelessly corrupt. These conclusions are not supported by Holly’s evidence, though flawed logic may lead her to believe that they are. Once again, false beliefs about logic do not change the character of one’s evidence and logical blindness is not a fast track to rationality.²⁵

What emerges from these reflections, therefore, is that in addition to factual beliefs stored in memory, a person’s CAB evidence will include the rules of logic and the relations of logical support that depend on them. Let us try to incorporate these additions into our account of CAB evidence:

CAB P belongs to S ’s CAB evidence at t_1 iff P belongs to S ’s CAD evidence at t_1 or S could access a memory of P at t_1 without relying on any additional sensory input or P is a rule of logic or P is logically supported by Q , where Q belongs to S ’s CAB evidence at t_1 .

This tentative modification raises a number of difficult questions: *Which* rules of logic are cognitively accessible to everyone?²⁶ How *strong* does the relation of logical support between P and Q need to be in order for P —the proposition supported—to be included in one’s evidence? If S believes Q , and Q supports P strongly and R weakly, is it *more rational* for S to believe P on the basis of Q than to believe R ? What about *other non-factual* propositions such as the basic truths of mathematics? These are important questions that will need to be answered in a complete statement of the theory (I do not have answers at this point). But such worries should not be allowed to obscure the crucial fact that we do hold people accountable not only for what they believe, but also for what follows as a matter of logic from what they believe. We say things like: “You should believe P , given that you believe Q ”; “It is not rational for you to believe P , given that you believe $Q, R, S \dots$ ”; “You believe P only because you are focusing on Q , and forgetting that you believe R , which clearly implies not- P .” If I am right that ADA is false—that we never have a duty to believe, disbelieve or suspend judgement about any proposition P —then these attributions of epistemic error (and the corresponding attributions of epistemic success) will need to be interpreted so as not to imply any duties to directly adopt any particular doxastic attitudes.

My suggestion is that epistemic oughts *as applied to particular propositions* are to be deflated. When we say that S epistemically *ought* (or *ought not*) to believe P given her evidence, we are appealing implicitly to the normativity of logic and we mean only that P is logically supported (or not logically supported) by S ’s CAB

evidence. To deflate an epistemic ought, then, is to treat sentences (1) and (2) below as equivalent:

- (1) *S* epistemically ought to believe *P* at t_1 on the basis of *S*'s evidence at t_1 .
- (2) *P* is logically supported by *S*'s evidence at t_1 .

It is important to recognise that the equivalence holds only if the rules of logic are a part of every person's CAB evidence. If we deny this, there will be situations where *P* is a logical consequence of some proposition that a person believes—(2) will be true—and yet it will not be true that the person ought to believe *P*—(1) will be false. So the deflationary theory that I propose depends on interpreting evidence widely. Another point is that Feldman's idea of a role ought might provide a useful way of understanding deflated epistemic oughts: If *P* is logically supported by *S*'s evidence at t_1 , then *S* is functioning well as a believer—*S* plays the role of believer successfully—if *S* believes *P* at t_1 on the basis of her evidence.²⁷

The case for deflating epistemic oughts as applied to particular propositions will need to be made out more fully in another place. Here, I just want to stress that my view is consistent with the existence of a more robust epistemic ought requiring action on our part. I think that this robust ought has its source in our capacity for rationality and is best understood as the general²⁸ epistemic duty to obtain proportion between our beliefs and our evidence. I therefore propose that we read Locke's principle in a way that is (1) consistent with basic doxastic involuntarism,²⁹ (2) imposes no duties to adopt any particular doxastic attitudes, and (3) recognises a general epistemic requirement to believe 'proportionably' with one's available evidence, where 'available evidence' includes the rules of logic broadly construed. I develop this idea in the final section and offer a limited defence.

3. The Duty to Obtain Proportion

The clear upshot of the previous sections is that if there is an epistemic duty to proportion our beliefs to our evidence, then it is one that we can fulfil only indirectly. My aim here is to adduce some reasons for thinking that we do have such a duty and to explain how and to what extent it can be met. I begin with another comment on the relation between deflated epistemic oughts and the normativity of logic.

When I say to you that *P* is logically supported by your evidence, I do not think that I need to give you any further argument in favour of your believing it. I just assume that as a rational being you have an interest in believing rationally, and I therefore hold you accountable to the standards of good thinking as determined by logic (again, 'logic' taken in a very wide sense). Evidentialists maintain that believing rationally is equivalent to believing in accordance with one's evidence. If I suppose

that you are a rational being, and I am an evidentialist, then I suppose that you should find believing in accordance with your evidence an epistemically desirable end, and so I suppose that you have a rational interest in believing *P*. Let's just suppose that evidentialism gets things right and see what follows from the falsity of ADA and the idea that judgements such as 'S ought to believe *P* given her evidence' in themselves express nothing more than the existence of a broadly logical relation between *S* and her evidence that we take to be binding on her.³⁰ (The assumption that evidentialism is true is one that many epistemologists reject. It is also one of the premises of my argument below.)

In general, it is very difficult to establish that we have duties of any sort. I think the best guide here (perhaps the only guide) will be our considered judgements about paradigm cases of duty or the absence of duty. A case of duty or its absence is paradigmatic, we'll say, if we can assume widespread agreement on the existence or nature of the duty in question, or widespread agreement on the absence of duty. Most people will agree, upon reflection, that we have a moral duty not to inflict gratuitous harm and that we have no duty, moral or otherwise, to look at the sun. These are paradigmatic cases, since our considered judgements about them will for the most part agree. I appeal to such judgements in making the following points.

First, I suggest that sentence (3) below does not follow from the conjunction of (1) and (2):

- (1) *M* is an action that would tend, either directly or indirectly, to bring about some end *E* that *S* ought to find desirable.
- (2) *S* has no duty to perform *M*.
- (3) *S* has no duty to bring about *E*.

Here is an interpretation on which (1) and (2) are arguably true while (3) is arguably false: *E* is the end of maintaining the health of one's body and *M* is the act of breathing. Most people's considered judgement would probably be that we have some sort of duty to maintain the health of our bodies and that rational beings desire this end. And it is certainly the case that breathing contributes to preserving the health of the body. But no one would say that we have a duty to breathe (since we cannot fail to breathe while living). The truth of (1) and (2) is also consistent with the falsity of (4):

- (4) *S* has no duty to take any action *N* that would tend, either directly or indirectly, to the accomplishment of *E*.

Let's say that *N* is the act of eating healthy foods. Eating healthy foods contributes to maintaining the health of the body and perhaps we have some sort of moral obligation to ourselves to eat at least some healthy foods. Whether or not this

is the case, the example makes clear that it is at least not implausible to maintain the falsity of (4) while asserting the truth of (1) and (2).

The next point is that role oughts imply duties only where we are dealing with beings capable of understanding their roles. There is some sense of 'ought' in which an American robin *ought* to build a nest, *ought* to eat berries and insects, *ought* to avoid chocolate (which is toxic to birds), and *ought* to migrate to a warmer climate in the winter. An American robin is functioning as it ought—fulfilling its role as this type of bird—insofar as it performs these actions. But we don't say that the bird has a duty to do any of these things in particular, nor a duty to maintain its health in general. By contrast, not only is a human functioning as she ought if she eats nutritious foods, avoids taking lethal drugs, gets a sufficient amount of exercise, and secures for herself as far as possible sanitary living conditions. Doing one's best to maintain a certain level of physical well-being is, in the view of many, a question of *duty* for humans. I do not assert that this is in fact one of our moral duties, but only that many people take it to be. And I agree that speaking of duty here is not entirely out of place as it would be in the case of the bird maintaining its health. The only relevant difference I can see between a bird fulfilling its role oughts and a human fulfilling hers is that the human is capable of understanding her roles. We can safely generalise the point and say that the capacity for reason is present in any case where we ascribe a duty.

Applying these points to belief, we get the following results. Our capacity to understand our role as believers makes it at least intelligible to speak in the strong terms of epistemic duty as opposed to the weaker terms of deflated epistemic (role) oughts. Furthermore, there is no valid inference from the absence of a duty to adopt any particular doxastic attitudes to the absence of epistemic duty altogether, nor to the absence of a duty to perform other actions that may lead indirectly to the accomplishment of some epistemically desirable end. The position I want to sketch now interprets our epistemic duty as the duty to obtain proportion between our beliefs and our evidence indirectly, and it does so in a way that makes the evidential availability of logic indispensable.

My view can be framed as an argument with ten premises. The first premise states that it is our epistemic duty to believe rationally. Though I can think of no non-questioning-begging argument to establish this, it seems to me entirely reasonable to think that it is the epistemic duty of a believer capable of reasoning to believe rationally. Perhaps there is an argument to be made from the fact that we do take the rules of good thinking to be binding on us. Or, if that is viciously circular, then perhaps we could say that this is just the most intelligible way to spell out what the normativity of logic—which we simply assume—implies in terms of action. We have ruled out duties to adopt particular doxastic attitudes—deflating epistemic oughts as they apply to particular propositions—and yet as rational beings we hold ourselves

accountable to logic in a general way: Logic tells us what we should or shouldn't believe on the basis of what we do believe. Since beliefs held logically are beliefs held rationally, I should, as a believer, believe rationally. But again, I suspect that any such argument begs the very question at issue, namely, whether or not we should believe rationally.

Perhaps the reason I can find no way to make the point without simply assuming it is that we are dealing with a conceptual truth. To classify someone as a rational being is just to say that we hold them accountable for their beliefs according to our shared standards of good thinking. A being capable of rational reflection on her beliefs is a being with an intellectual obligation to believe rationally, just as a being capable of moral reflection on her actions is a being with a moral obligation to act morally.

Looking at it from a different angle might help: Those who deny that there is such a thing as morally right or wrong action will not classify anyone as capable of *moral* reflection, since moral reflection is a matter of reflecting on what one morally should or shouldn't do. What others call 'moral reflection' they might describe as prudential reflection misconstrued, or as just plain nonsense. But they won't allow that anyone is actually engaged in genuinely moral reflection. Similarly, it wouldn't make any sense for me to classify you as a rational being if I didn't think that there were anything epistemically objectionable in your forming all of your beliefs irrationally. In short, if I regard you as moral I regard you as somehow answerable to standards of morality; if I regard you as rational I regard you as somehow answerable to standards of rationality. So to say that a rational being should believe rationally is to say the same thing twice. Though these reflections hardly establish the first premise of my argument, they at least suggest a standpoint from which a plausible defence could be mounted.³¹

The second premise says that a person believes rationally when their beliefs are justified. This is just a matter of definition. What I mean by 'believing rationally' is believing *with justification* rather than *without* it. The difficult question is what epistemic justification itself consists in, and whether the factors conferring justification on beliefs are all internal to the believer's cognitive states. But this is not an issue that needs to be settled here, since whatever theory one adopts, there doesn't seem to be any good reason to distinguish between beliefs held rationally and beliefs held justifiably. A reliabilist, for example, could argue that a belief is justified—that is, held rationally—when it issues from a reliable belief-forming process. And a coherentist could argue that a belief is justified—that is, held rationally—when it belongs to a coherent set of mutually supporting beliefs.³² The definition stated in premise (2) is thus consistent with any of the standard theories of epistemic justification, but I spell it out and resist collapsing it into premise (1) for the sake of clarity.

Premise (3) is contentious since it asserts without argument that evidentialism

is the correct theory of epistemic justification. In my view, a person's beliefs are justified if they are proportionate with their evidence. If S 's evidence supports P , then the justified (or fitting) doxastic attitude for S to adopt toward P is belief. If S 's evidence supports $\sim P$, then the justified doxastic attitude for S to adopt toward P is disbelief. If S 's evidence supports neither P nor $\sim P$, then the justified doxastic attitude for S to adopt toward P is suspension of judgement.³³ In short, premise (3) says that a person's beliefs are justified when they are proportionate with their evidence. A defence of evidentialism is clearly beyond the scope of this paper, so I merely flag this as work that needs to be done in order for my argument to succeed.

The fourth premise follows straightforwardly from the first three: (4) Our epistemic duty is the duty to proportion our beliefs to our evidence. The question of how we go about fulfilling this duty is complicated by the fact that we have, according to the previous argument, no duty to adopt any particular doxastic attitudes. The duty to obtain proportion between belief and evidence would be met directly for a person S if, for any proposition P that S considers, S were to adopt the fitting or justified doxastic attitude towards P immediately upon considering P (not spontaneously, of course, but by exercising basic doxastic voluntary control, since S can have a duty to adopt a particular doxastic attitude toward P only if S can fail to have that doxastic attitude toward P). But the falsity of ADA implies that this is not what is epistemically required of us: The concept of duty does not apply directly to belief, disbelief or suspension of judgement—this is premise (5). So it is our epistemic duty to obtain the desired proportion indirectly—this is premise (6).

If our epistemic duty to obtain proportion can be met indirectly, then there must be other things we can do to bring our beliefs into line with our evidence, things which it will therefore be our epistemic duty to do. Feldman thinks that we exercise non-basic immediate voluntary control over some of our beliefs, specifically, over beliefs about states of the world which are themselves within our voluntary control, provided that our beliefs track those states accurately (Feldman 2004a: 171). To use Feldman's example, suppose that I am sitting in my office in the dark thinking about the fact that the lights are off. I cannot through a simple act of the will give myself the belief that the lights are on. But I do have non-basic immediate voluntary control over whether the lights in my office are on or off. If I walk over and flip the switch, the lights will come on. And given that my belief tracks that state of the world accurately, I will immediately form the belief that the lights in my office are on. This shows that I have non-basic voluntary control over some of my beliefs: There are things I can do to bring about certain states of the world, and where my beliefs are responsive to those states I can give myself the corresponding beliefs. The same reasoning will show that I have long range voluntary control over some of my beliefs. I can give myself the belief that a wardrobe exists in my office by building one there.

I think that the sort of voluntary control that Feldman identifies here can be seen as a metaphor for another type of control to which our beliefs are subject. This other type of control is both immediate non-basic and long range,³⁴ but it is relevant to the epistemic goal of proportioning our beliefs to our evidence in a way that the immediate non-basic control described by Feldman is not. To see this, let's first elaborate on Feldman's example. True, I can move my arm in a certain way to give myself the belief that the lights are on (by flicking the switch). I could also induce the belief that I am holding a pencil by picking one up (since my belief will track that state of the world equally well). But I don't actually need to move any limbs to give myself beliefs about certain states of the world. Suppose I want to induce the belief that my colleague, Jim, is in my office. To do that, I could just shout, "Jim, would you please step into my office for a minute?" If Jim complies, then I will have given myself the belief that Jim is in my office (for when I see him walk in, I will instantly form the belief that he is in my office.) Now suppose that Jim walks in and, gesturing towards the window behind me, says, "Look at that bird in the tree out there!" If I have no reason to doubt Jim's veracity, then, on the basis of his testimony, I will form the belief that there is a bird in the tree behind me. Here I have less control. Finally, suppose that I want to form a belief about the particular colour of the bird in the tree behind me. Assuming that there are several different species in the neighbourhood, I am powerless to give myself the belief that the bird is of one colour rather than another by simply willing it. But it is in my power to turn my head and train my sight on the tree behind me. This would presumably cause me to form a belief about the colour of this individual bird. Again, I have triggered a belief by exercising immediate non-basic voluntary control over states of the world, in this case by controlling the direction in which my head is facing and my eyes are focused.

What all of these examples have in common is the input of new sensory data: In each case my belief is responsive to information that was not previously contained in my CAB evidence and thus my total available evidence has grown. Feldman's example therefore shows that we can exercise non-basic immediate voluntary control over some of our beliefs by manipulating the world in such a way as to give ourselves new sensations. This doesn't tell us anything about rationality, however, since believing rationally is not a matter of getting more evidence, but of understanding the evidence that we do have.³⁵ This is precisely how I think we should interpret Locke's principle of proportionality: The epistemic duty to proportion our beliefs to our evidence is the duty to understand the evidence we already possess. We can accomplish this effectively only by deliberating on our evidence, that is to say, by thinking carefully about the logical relations that hold between our evidence and the propositions we believe or consider. The fitting doxastic attitudes will follow automatically.

Return for a moment to Jim's testimony about a bird sitting in the tree behind me. Suppose that I know Jim to be an inveterate liar and a prankster who has deceived me many times before. Let's say that I am also a bird enthusiast who wants more than anything to see birds in trees. If in my zeal for bird sightings I block my knowledge of Jim's mendacity from conscious thought—I don't allow myself to cognitively access all of the pertinent information that I possess—and I take his testimony at face value, then I am decidedly not proportioning my belief to my available evidence. The rational response to Jim's assertion, given what I know, would be for me to suspend judgement until I can acquire more information. Without even turning my head, I may be able to give myself the fitting doxastic attitude (suspension of judgement) by deliberating on my available evidence. Reflection on what I know and how it relates to Jim's assertion will guide my belief to the propositions my evidence supports.³⁶

The act of deliberating on our evidence is merely analogous to the acts by which we control states of the world to which our beliefs are responsive. When I reflect on my knowledge of Jim's character, for instance, and think about how this knowledge relates to Jim's recent assertion, I am in the process of coming to understand my existing evidence. By contrast, when I turn and look out the window at the bird in the tree, I am acquiring new evidence. To generalise the point, whereas controlling states of the world leads to the acquiring of new evidence, critical reflection on evidence already possessed leads to believing rationally (or 'proportionably', as Locke would say).³⁷ Once we understand our evidence, the justified doxastic attitudes follow by psychological necessity and due proportion is achieved.³⁸ Accordingly, premise (7) says that we can attain proportion between our beliefs and our evidence indirectly only by deliberating on our evidence.³⁹

Premise (8) combines (6) and (7): Our epistemic duty can be met only by deliberating on our evidence. But it is not possible for a person to think constantly about the totality of logical relations that hold between her evidence and every proposition that she believes or considers. This is premise (9). Any real person has more beliefs than she can possibly contemplate at one time, let alone constantly. The constraint here is not merely psychological. Every bit of sensory experience I have supports an infinite number of propositions, and any proposition that I believe has an infinite number of logical consequences, so it is in principle impossible for me to reflect on all of them. From the "infinite justificational fecundity" of propositional knowledge and sensory experience, Mark Nelson infers that we have no positive epistemic duties (Nelson 2010).

But this inference is not obviously valid. As we saw above, the fact that a person cannot fulfil an obligation does not mean that she does not have the obligation.⁴⁰ I have argued that we do have an intellectual obligation to proportion our beliefs to our evidence (and, for different reasons than Nelson offers, that we have no intellectual obligation to believe any particular propositions). Perhaps the justificational

fecundity of sensory and propositional evidence implies rather that the obligation to obtain proportion can be met only imperfectly. This idea fits well with premise (8), which provides guidance in relatively general terms ('Deliberate on your evidence'), but does not rule out any specific kinds of actions ('Do not fall for red herrings', 'Never accept unreliable testimony', 'Do not infer $\sim Q$ from $\sim P$ and $P \supset Q$ '). My recommendation is that we interpret the intellectual requirement expressed under (8) as the epistemic equivalent of an imperfect moral duty in Kant's sense.⁴¹

On one reading of Kant, imperfect (moral) duties have the following features:

1. Imperfect duties prescribe general guiding principles for promoting certain ends (as opposed to proscribing specific kinds of actions).
2. Imperfect duties do not prescribe particular actions.
3. It is up to the agent to determine when and to what extent an imperfect duty is to be satisfied. (Their structure is: "One ought to do (or avoid) x sometimes, to some extent" (Hill 1992: 149)).
4. Judgement is required to determine when and to what extent particular actions are to be performed in order to satisfy the general end prescribed by an imperfect duty.
5. Imperfect duties assign no determinate upper limit to what must be done (though there is a minimal threshold below which the imperfect duty is not met).
6. Satisfying an imperfect duty by performing a particular action on some occasion is always praiseworthy (whereas violating a perfect duty on any occasion is always blameworthy).⁴²

If we interpret Locke's proportionality principle as an imperfect epistemic duty, as I think we should, then the following statements will be true: We must deliberate on our evidence sometimes and to some extent. There are no particular errors of reasoning that it is our epistemic duty to avoid always. We have freedom to choose to deliberate on our evidence, or not, on a given occasion, provided that we are prepared to deliberate carefully on some other occasions.⁴³ We must use our judgement to determine how much deliberation is required, when it is required, and on which propositions. There is no determinate limit to how much deliberating we should do (though there is an unspecified minimum).⁴⁴ It is always epistemically praiseworthy for us to deliberate on our evidence, though failing to deliberate on some occasions is not epistemically blameworthy.⁴⁵ I think it is a strength of the theory I am advocating that it has these results.⁴⁶

The analogy between imperfect moral and epistemic duties can be pushed a little further. A paradigmatic example of an imperfect moral duty for Kant is the duty to help those in need. Though the duty to help those in need is absolutely binding on

me, because it is imperfect there is some latitude in how I go about fulfilling it. A final example will make this clear. Imagine that there is a student in South Africa named Sifiso who has real talent for scientific work, but lacks the funding and resources to continue his studies beyond the undergraduate level. Suppose as well that I have no knowledge of Sifiso's existence. Clearly, I have no moral duty to assist Sifiso financially (even if I knew of his existence, I would still have no duty to help him rather than, say, an American or British student in similar circumstances, or to help someone in entirely different circumstances). But helping Sifiso would certainly be a morally good thing for me to do. Not only would it be morally desirable for me to help Sifiso (as it would be epistemically desirable for me to believe P on the basis of my evidence, if P is a remote logical consequence of my beliefs). Supporting Sifiso would also count towards the fulfilment of my imperfect moral duty to help those in need. Similarly, discovering through deliberation that P is a logical consequence of my beliefs—and believing P (automatically) as a result—would contribute towards the fulfilment of my imperfect epistemic duty to proportion my beliefs to my evidence, though I had no duty to figure out that P specifically (rather than Q , R or S) is thus entailed. So although the concept of duty does not apply directly to belief, the act of deliberating on my evidence influences belief and contributes positively to the general epistemic end that I am duty-bound to pursue.⁴⁷

This gives us premise (10) of the argument, which can now be summarised as follows:

The Argument for PBE as an Imperfect Epistemic Duty

- (1) Our epistemic duty is the duty to believe rationally.
- (2) A person believes rationally when their beliefs are justified.
- (3) A person's beliefs are justified when they are proportionate with their evidence.
- (4) Our epistemic duty is the duty to proportion our beliefs to our evidence (PBE is true). (1), (2), (3)
- (5) We have no duty to adopt any particular doxastic attitudes (ADA is false).
- (6) Our epistemic duty is the duty to proportion our beliefs to our evidence indirectly. (4), (5)
- (7) A person can attain proportion between their beliefs and their evidence indirectly only by thinking carefully about the logical relations that hold between their evidence and the propositions they believe or consider.
- (8) Our epistemic duty can be met only by thinking carefully about the logical relations that hold between our evidence and the propositions that we believe or consider. (6), (7)

- (9) It is not possible to think constantly about the logical relations that hold between our evidence and every proposition we believe or consider.
- (10) Our epistemic duty can only be met imperfectly. (8), (9)
- (11) Our epistemic duty can be met only imperfectly by thinking carefully about the logical relations that hold between our evidence and the propositions that we believe or consider. (8), (10)

The basic idea behind the entire argument is that rationality consists not in acquiring new evidence, but in understanding the evidence one already has. Evaluation of doxastic attitudes for rationality makes sense only if irrationality is possible (S believes P rationally at t_1 only if it would be irrational for S to believe $\sim P$ at t_1 .) And it makes sense to say that a doxastic attitude is held irrationally—that is, to say that a person has failed to understand their evidence—only if the rules of good thinking in general are regarded not as data to be acquired, but as available upon reflection and thus as cognitively accessible to anyone whom we call rational. If it is proper for a rational being to believe rationally, and if rationality is achieved through deliberation on evidence possessed, then, given the infinite justificational fecundity of evidence, the epistemic duty to deliberate is best construed as imperfect—absolutely binding, flexible in application, and prescribing perfect proportion between belief and evidence as an unattainable ideal.

Much work remains to be done to fill in the details of the theory I have advanced and to offer additional support for the underlying assumptions. For now, I hope that the power of the theory to solve difficult problems will be taken as a preliminary argument in its favour.

References

- Adler, J. 2002. *Belief's Own Ethics*. Cambridge: MIT Press.
- Alston, W. 1988. The Deontological Conception of Epistemic Justification. *Philosophical Perspectives* 2: 257–99.
- Butchvarov, P. 2008. Epistemology Dehumanized. In: Q. Smith (ed.) *Epistemology: New Essays*. New York: Oxford University Press.
- Conee, E. and Feldman, R. 2004. Evidentialism. In: E. Conee and R. Feldman (eds.) *Evidentialism: Essays in Epistemology*. New York: Oxford University Press.
- . 2008. Evidence. In: Q. Smith (ed.) *Epistemology: New Essays*. New York: Oxford University Press.
- Davidson, D. 1989. A Coherence Theory of Truth and Knowledge. In: E. LePore (ed.) *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*. New York: Blackwell.
- Feldman, R. 2002. Epistemological Duties. In: P. Moser (ed.) *The Oxford Handbook of Epistemology*. New York: Oxford University Press.

- . 2003. *Epistemology*. Upper Saddle River, NJ: Prentice Hall.
- . 2004a. The Ethics of Belief. In: E. Conee and R. Feldman (eds.) *Evidentialism: Essays in Epistemology*. New York: Oxford University Press.
- . 2004b. Having Evidence. In: E. Conee and R. Feldman (eds.) *Evidentialism: Essays in Epistemology*. New York: Oxford University Press.
- . 2004c. Authoritarian Epistemology. In: E. Conee and R. Feldman (eds.) *Evidentialism: Essays in Epistemology*. New York: Oxford University Press.
- Ginet, C. 2001. Deciding to Believe. In: M. Steup (ed.) *Knowledge, Truth and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. New York: Oxford University Press.
- Hill, T. 1992. *Dignity and Practical Reason in Kant's Moral Theory*. Ithaca: Cornell University Press.
- Hume, D. 1990. *Enquiries Concerning Human Understanding and Concerning the Principles of Morals*, ed. L. A. Selby-Bigge, 3rd edn, rev. P. H. Nidditch. Oxford: Oxford University Press.
- Kant, I. 1996. *Practical Philosophy*, trans. M. J. Gregor and A. Wood. New York: Cambridge University Press.
- Katzoff, C. 2000. Counter-Evidence and the Duty to Critically Reflect. *Analysis* 60: 89–96.
- Kim, K. 1994. The Deontological Conception of Epistemic Justification and Doxastic Voluntarism. *Analysis* 54: 282–84.
- Locke, J. 1975. *An Essay Concerning Human Understanding*, ed. P. H. Nidditch. New York: Oxford University Press.
- Nelson, M. 2010. We Have No Positive Epistemic Duties. *Mind* 119: 83–102.
- Reynolds, S. Doxastic Voluntarism and the Function of Epistemic Evaluations. *Erkenntnis* 75: 19–35.
- Steup, M. 1996. *An Introduction to Contemporary Epistemology*. Upper Saddle River, NJ: Prentice Hall.
- Tidman, P. 1996. Critical Reflection: An Alleged Epistemic Duty. *Analysis* 56: 268–76.
- Wolterstorff, N. 1996. *John Locke and the Ethics of Belief*. New York: Cambridge University Press.

SCOTT STAPLEFORD
 Department of Philosophy
 Faculty of Humanities
 St. Thomas University
 Fredericton, New Brunswick E3B 5G3
 CANADA
 sstapleford@stu.ca

Resumo. Este artigo defende a tese de que nosso dever epistêmico é harmonizar nossas crenças à evidência que possuímos. Propõe-se uma visão abrangente da evidência possuída que faz sentido de nossas intuições a respeito de quando é correto dizer que uma pessoa deve acreditar em alguma proposição *P*. Uma segunda tese é que não temos nenhum dever epistêmico de adotar quaisquer atitudes doxásticas particulares. A tensão aparente entre as duas teses é resolvida aplicando-se indiretamente o conceito de dever à crença.

Palavras-chave: Deveres epistêmicos; étic da crença; voluntarismo doxástico; evidência; evidencialismo; R. Feldman.

Notes

¹ I attribute this view to Feldman on the basis of an argument that occurs in Section II B of “The Ethics of Belief” (Feldman 2004a: 181).

² The term ‘principle of proportionality’ comes from Nicholas Wolterstorff (see Wolterstorff 1996: 79).

³ For convenience, I will sometimes speak simply of duties to believe, but whatever I say should always be taken to hold for the doxastic attitudes of disbelief and suspension of judgement as well.

⁴ An interesting example is the ancient challenge to believe that ‘The number of stars is even.’ (See Adler 2002: 25). It will be useful in what follows to call the thesis that we do exercise basic voluntary control over our beliefs *basic doxastic voluntarism*.

⁵ Here I follow Feldman’s summary of the argument very closely (Feldman 2004a: 168).

⁶ Carl Ginet (2001) takes a hard line in defending the voluntariness of belief. Less extreme is the view of Matthias Steup, who recommends that deontologists respond to this argument by distinguishing between the *soft* and *hard* involuntariness of belief. (See Steup 1996: 78–9). Premise (2) will come out true or false depending on the interpretation of ‘soft’. Reynolds (2011) argues that we do have control over our beliefs, although we cannot simply believe at will.

⁷ For reasons that will become clear in Section 2, I speak interchangeably of logical support and evidential support, where ‘logic’ is construed widely to include all forms of good reasoning.

⁸ I call this a ‘deflationary’ sense of ‘ought’. See the end of Section 2 for discussion. Note that not all cases where the word ‘ought’ applies are cases where the word ‘duty’ applies (though the reverse does hold). Your heart *ought* to pump blood just insofar as it is the kind of organ that it is (this may be a role ought). But your heart does not have a *duty* to pump blood.

⁹ In Section 3, I will argue that the concept of duty applies to belief in a less direct fashion.

¹⁰ Let’s assume for the sake of argument that the supporting relation is strong enough to make the proposition at least highly probable. This will allow us to bracket the difficult question of how to align degrees of evidential support with appropriate degrees of belief.

¹¹ Perhaps the connection between apprehension of evidence for *P* and belief that *P* is stronger than psychological. Charlotte Katzoff claims that there is a logical connection between *S*’s view of her evidence and the beliefs that *S* attributes to herself (see Katzoff 2000: 92–3). My argument requires only that *S* will believe *P* without fail if *S* recognises that her total evidence (strongly) supports *P*, but this is consistent with there being in addition a conceptual link between our perception of our evidence and the doxastic attitudes we ascribe to ourselves. But compare note 14 below.

¹² Assume whatever background information is needed in order to rule out extraordinary defeaters of this evidence (I am not on perception-altering medication, I have no reason to believe that someone is trying to trick me, and so on).

¹³ Notice that this assumption alone implies nothing about the epistemic evaluation of belief.

¹⁴ One referee for *Principia* asks whether I want this to be an empirical claim or an analytic claim. I think it is an empirical claim based on an induction from introspection. We are all aware of our inability to control our beliefs directly when faced with overwhelming evidence: I cannot fail to believe that my girlfriend is now sitting on a chair beside me, that the tea I just sipped is hot or, when I think about it, that $\sqrt{9}$ is 3. By contrast, I don't feel compelled either way about the proposition that my girlfriend is wearing pink socks (I can't see her feet), the proposition that the caffeine content of my tea is higher than it was in the last cup, or (without attempting a proof) the proposition that $(\exists x)(P \supset Qx)$ implies $(P \supset (\exists x)Qx)$. I suppose that in principle systematic empirical studies could be done to test the claim, but they would be exceedingly difficult to set up. A researcher would have to figure out both what a subject's total body of evidence comprises and what the evidential relations are between that body of evidence and various propositions considered. The researcher would then have to determine whether or not the subject believes some logical consequence of her evidence and whether or not she believes it on the basis of her evidence. Also, the research would be unmotivated, since the question doesn't seem very urgent.

¹⁵ I include this example from Feldman because it illustrates the point that we can have duties to do things even when fulfilling them is beyond our voluntary control (Feldman 2004a: 173).

¹⁶ These categories are not mutually exclusive: If *S* has cognitively accessed *P*, then *P* is *ipso facto* cognitively accessible to *S*.

¹⁷ Feldman considers and rejects a similar account of 'availability' (*vide* Feldman 2004b: 231). Some philosophers believe that only propositions or beliefs can be potential justifiers, and that perceptual, memorial and other experiences do not count as evidence (see, for instance, Davidson 1989). I reject this view, but it will be best to focus on propositions here for the sake of clarity.

¹⁸ Feldman argues that a better account of evidence would construe availability in terms of what is easily accessible or easily retrievable (Feldman 2004b: 231). In the end he rejects this account as well because of obscurity in the notion of easy accessibility. But the fact that the notion of easy accessibility is obscure does not mean that it is the wrong notion.

¹⁹ By 'non-factual propositions' I mean, roughly, non-empirical propositions.

²⁰ For the purposes of this paper, I want to interpret 'logic' in the widest sense possible. Principles of deduction, induction, probability and abduction are all 'logical' principles in this sense, as are any rules of good thinking that do not fit neatly under one of these headings, such as principles expressing relations of non-formal entailment (on the latter, and the related idea of epistemology-as-logic, see Butchvarov 2008: 318ff.) In other words, I want to take 'logical support' to be equivalent to 'evidential support'. I also do not want to rule out perceptual, memorial or other experiences as potential justifiers. I think that Conee and Feldman are right to classify fundamental epistemic principles as principles of best explanation: "Perceptual experiences can contribute toward the justification of propositions about the world when the propositions are part of the best explanation of those experiences that is available to the person. Similarly, the truth of the contents of a memory experience may be part of the best explanation of the experience itself" (Conee and Feldman 2008: 97–8). Since I include principles of best explanation (abductive principles) within logic, these epistemic principles are broadly logical principles. Your perceptual experiences can therefore provide logical support for certain propositions if those propositions are a part of the best

explanation of the experiences. But I want to leave non-propositional, experiential justifiers aside as an unnecessary complication for present purposes.

²¹ By ‘propositions expressing relations of logical support’ I mean such things as ‘*P* implies *Q*’, ‘*P* makes *Q* probable’, ‘*P* inductively supports *Q*’, and so on.

²² The reason that Feldman draws the opposite conclusion from the example is that he has a more generous view of what it is to believe rationally: *S* believes *P* rationally to the extent that *P* is evidentially supported by the propositions that *S* is actually thinking of, even if *S* has non-accessed counter-evidence to *P* stored in memory. This more generous view is of course grounded in his more restricted view of what it means to have something available as evidence: *S* has *P* available as evidence if and only if *S* is actually thinking of *P* (if and only if *P* belongs to *S*’s CAD evidence, in my terminology). Feldman may now be backing away from this extreme view of what counts as the evidence a person has (see Conee and Feldman 2008: 89). He admits, though, that he vacillates on the point and that he and Conee are not always in agreement.

²³ Note that this does not imply that the subject has a *duty* to adopt any particular doxastic attitude with respect to *P* (this follows from the argument of Section 1). To say that *S* ought not believe *P* means only that *P* is not supported by *S*’s evidence. So epistemic oughts, though they may imply certain indirect duties with respect to belief—as I will argue in Section 3—do not imply any duties to believe, disbelieve or suspend judgement about any particular propositions. It should be obvious that my view relies on the truth of evidentialism, a theory of epistemic justification that I assume here without argument. This is of course a limitation of my position.

²⁴ An argument against my position could be constructed from an example that occurs in Feldman 2004c: 114–15. I hope to develop this argument and respond to it in a separate paper.

²⁵ A similar claim can be made regarding stored beliefs about factual matters: Failing to take note of evidentially relevant memories may lead to the holding of beliefs contrary to one’s available evidence and thus to the epistemically undesirable state of believing something irrationally. For example, suppose that Holly fails to keep an appointment with her partner, Jordan, who, we will say, has countless proofs in memory of Holly’s fidelity. If Jordan allows anger and jealousy to cloud his judgement to such an extent that he comes to believe that Holly must be having an affair, he is believing something contrary to his evidence and thus irrationally. His knowledge of Holly’s good character, built upon the many experiences he has of her, should trump the suspicion raised by one missed appointment. The fact that emotion is currently blocking Jordan’s memories of Holly’s past behaviour does not make it rational for him to conclude that she is being unfaithful.

²⁶ I am open to the possibility that very few rules of logic are cognitively accessible to everyone—only the ‘easy’ ones, as one referee put it—and that the more exotic theorems of the various non-classical systems are not. But if this is so, my theory will make a mess of evidential relations. If axiom S4 is *in principle* unavailable to a subject *S* and *S* believes $\Box P$, for instance, then *S*’s body of evidence will not logically support $\Box\Box P$ (assuming *S* believes nothing else that would enable the inference to $\Box\Box P$ using rules that are available to *S*). This follows from my assumption that a rule of logic is valid only if it is in principle accessible through reflection. Exotic theorems are valid, in other words, only if they are reasonable. But I have no clear view on which rules of logic are valid. What is clear, I think, is

that the evidentialist must hold that the propositions it is rational for a subject to believe on the basis of her evidence are the propositions that are logically supported by her evidence. The evidentialist can make this claim without in fact being able to say which propositions are logically supported by which bodies of evidence, since there is no requirement for an evidentialist to be logically omniscient. Perhaps my theory of epistemic obligation could be strengthened by coupling it with logical pluralism, the view that there is more than one right logical system and that the system that it is appropriate to reason with depends on one's epistemic circumstances. In that case, whether or not *P* is supported by *S*'s evidence depends in part on which logical system it is appropriate for *S* to reason with in her current epistemic circumstances.

²⁷ The qualification 'on the basis of her evidence' is important. If *S* believes *P* not because *P* is logically supported by her evidence, but as a result of wishful thinking, then *S* is not functioning well as a believer. An interesting objection to my position might be developed by introducing a distinction between current-state justification and all-evidence-considered justification. If *P* is a logical consequence of *S*'s evidence and *S* fails to see the connection between her evidence and *P*, then *S* will be all-evidence-considered justified in believing *P* (since *P* is supported by *S*'s total available evidence), but not current-state justified in believing *P* (since *P* is not supported by *S*'s CAD evidence). So *S* will be both justified and not justified in believing *P*. While I see the appeal of introducing such a distinction, I resist doing so for the reasons laid out above. To advert to earlier examples, Holly is decidedly not justified in believing that all horses are white on the basis of her evidence, and Jordan is not justified in believing that Holly is cheating on him on the basis of his. I avoid the objection, in other words, by rejecting the category of current-state justification. Many people believe many things irrationally by failing to consider the evidence that is cognitively available to them. Allowing the category of current-state justification will force us to call numerous doxastic attitudes rational or justified which are clearly not. In fact, it is difficult to see how *any* doxastic attitudes could be ruled out as current-state irrational on such a view. If I genuinely believe *P* as a result of wishful thinking, then I am currently taking wishful thinking to be a good reason for believing something and not currently recognising anything else as a sufficiently strong reason for disbelieving *P*. So *P* will be current-state justified for me. But this is not an example of justified or rationally held belief, so the current-state conception which says that it is must go.

²⁸ The duty is general in that it does not apply to any particular propositions or doxastic attitudes.

²⁹ Doxastic involuntarism is just the negation of doxastic voluntarism. It says that belief is *not* within our voluntary control.

³⁰ In other words, though such judgements do not in themselves require any action on our part (as the duty to obtain proportion between our beliefs and our evidence does), they do tell us what it would be epistemically desirable for us to believe on the basis of our evidence.

³¹ Feldman tries to show that epistemological duties are duties we must carry out if we are to have justified beliefs (see Feldman 2002: 373–80). This argument, if successful, would establish the first and second premises of my argument.

³² Some epistemologists prefer to speak of warrant rather than justification, and they equate rational belief with warranted belief. 'Justified' can be replaced with 'warranted' in premise (2) as a terminological point, but then premise (3) will have to be rejected on any theory

that interprets ‘warrant’ in non-evidentialist terms.

³³ For simplicity I abstract again from degrees of evidential support and proportionate degrees of belief (strong belief for strong evidence, weak belief for weak evidence, and so on).

³⁴ Feldman is right to point out that the boundary between immediate non-basic and long range voluntary control is imprecise (Feldman 2004a: 170).

³⁵ I think Feldman is right to reject what he calls the ‘Get the Evidence Principle’, according to which there are situations in which we have an epistemic duty to acquire additional evidence. Though getting more evidence may be a generally desirable thing, believing rationally does not depend on it. As Feldman says: “Even when one should get more evidence, the thing to do at any given time [according to evidentialism] is to be guided by the evidence one does have” (Feldman 2003: 48).

³⁶ Feldman would no doubt regard the act of coming to understand one’s evidence as equivalent to the acquiring of new evidence, but I think that this is a mistake, if the argument of Section 2 is correct. If I believe that $P \& Q$, then Q is a part of my CAB evidence, even if I am not thinking of the conjunction elimination rule.

³⁷ My point can also be made negatively. Irrationality does not consist in not having enough evidence, but in failing to understand the evidence one does have. Since it is our epistemic duty to believe rationally, it is our duty to come to understand our evidence. The only way to do that effectively is by deliberating on our available evidence.

³⁸ One might argue that this manner of guiding belief falls short of voluntary control. For an action to be within our voluntary control an intention to perform that action must be formed. When we guide our beliefs by deliberating on our evidence, no intention to form any particular belief arises. Therefore, guiding belief through deliberation on evidence does not amount to voluntary control. (Alston and Feldman argue along these lines. See Alston 1988: 271 and Feldman 2004a: 172–3.) I have no reason to oppose this argument, since the coherence of my own position does not depend on the voluntariness of belief. As we saw in Section 1, belief need not be voluntary for it to be open to epistemic evaluation, and my view does not require agents to form intentions to take up particular doxastic attitudes. The essential thing for me is that it is within our power to deliberate on our evidence. If we deliberate well, fitting doxastic attitudes will follow and our beliefs will be proportionate with our evidence. This is a way to affect belief indirectly, and since it leads to rationality it is our epistemic duty. But I can easily give up the term ‘voluntary’ as applied to belief itself.

³⁹ Perhaps we can also influence our beliefs appropriately by exerting very long range voluntary control, say by enrolling in critical reasoning courses. But I doubt whether this sort of activity would count towards satisfying our epistemic duty. It may be merely desirable (though it would most likely improve our proficiency in deliberating on our evidence).

⁴⁰ I am simplifying Nelson’s argument for the sake of clarity.

⁴¹ The distinction between perfect and imperfect duties is introduced at *Groundwork for the Metaphysics of Morals*, 4:222-23 (Kant 1996: 73).

⁴² I draw on Thomas Hill’s account of imperfect moral duties for some of the details (see Hill 1992: 148–60).

⁴³ I am paraphrasing Hill’s wording here and translating it into epistemic terms (*cf.* Hill 1992: 155).

⁴⁴ Though we probably cannot fail to deliberate on our evidence at least sometimes, we can

fail to perform this imperfect duty, since there is a minimal threshold of deliberation below which we fail to do what is epistemically required of us—we can fail to reflect to an adequate degree or with enough frequency or sincerity. This is a distinctive feature of imperfect duties in general: We can do too little to meet them and we can always do more to fulfil them.

⁴⁵ I am not suggesting that it is always all-things-considered best to deliberate on one's evidence. There may be practical reasons for not doing so.

⁴⁶ One benefit of this account is that it will take care of some of the worries raised by Tidman (1996) and Katzoff (2000) regarding Kim's (1994) attempt to defend deontological theories of epistemic justification by appeal to the ways we can control belief indirectly. This is something to be pursued elsewhere.

⁴⁷ Obviously this notion of an imperfect epistemic duty will need to be developed in much greater detail. All I am doing here is introducing the theory.