EVIDENCE AND FALLIBILITY

JOSHUA DIPAOLO
joshuadipaolo@gmail.com

ABSTRACT

The "Evidentialist Dictum" says we must believe what our evidence supports, and
the "Fallibility Norm" says we must take our fallibility into account when man-
aging our beliefs. This paper presents a problem for the Evidentialist Dictum
based in the Fallibility Norm and a particular conception of evidential support.
It then addresses two novel Evidentialist responses to this problem. The first
response solves the problem by claiming that fallibility information causes "evi-
dence-loss." In addition to solving the problem, this response appears to explain
what's wrong with certain illegitimate dismissals of misleading evidence.
However, this explanation opens it up to objections. Next, I consider and pose
challenges to an Evidentialist strategy that attempts to solve the problem by
converting accounts of fallibility's epistemic significance for rational belief into
principles of evidential support. I conclude by sketching a solution that allows
us to capture what's true in the Evidentialist Dictum and the Fallibility Norm.

A wise man … proportions his belief to the evidence.

– Hume, *An Enquiry Concerning Human Understanding*

In general, there is a degree of doubt, and caution, and modesty, which, in all kinds of scrutiny and
decision, ought for ever to accompany a just reasoner.

– Hume, *An Enquiry Concerning Human Understanding*[1]

1. INTRODUCTION

This paper presents a problem for the Evidentialist Dictum based in the Fallibility Norm.
According to the former, we must believe what our evidence supports, while according to
the latter, we must take our fallibility into account when managing our beliefs. Both
claims are plausible, but they appear to conflict.[2]

To illustrate, consider the peer disagreement debate. Conciliationists insist, on the basis
of fallibility considerations, that agents should often reduce their confidence when con-
fronted with a disagreeing peer.[3] Others accuse conciliationists of inappropriately

---

1   Hume (1999: 170, 208).
2   Work addressing this conflict includes: Feldman (2005, 2009), Kelly (2005, 2010), Christensen (2007,
    2010, 2011, 2013), Elga (2007), Schechter (2013), Weatherson (2013, ms), Horowitz (2014),
    Lasonen-Aarnio (2014), Schoenfield (2015), Sliwa and Horowitz (2015), Titelbaum (2015), and
    DiPaolo (2016).
3   See Christensen (2007, 2013), Elga (2007), Feldman (2007), and Kornblith (2010).

requiring agents to *ignore evidence*.[4] Even some conciliationists believe epistemic modesty has its limits: the evidence will *always* favor conciliationism, and no amount of disagreement can make ignoring *that* evidence reasonable.[5]

The conflict isn't limited to disagreement. The Evidentialist Dictum and the Fallibility Norm often conflict because taking account of our vulnerability to error when assessing evidence often requires us not to fully trust our assessments. In turn, this can require us to believe against our evidence. For this reason, the Fallibility Norm cannot be properly accommodated by any view that implies epistemic rationality consists solely in heeding the Evidentialist Dictum. This is a problem for those views.

I address and ultimately challenge two defenses of the Evidentialist Dictum. The first claims that fallibility considerations take effect by causing "evidence-loss" (§§2–3). This approach has implausible implications about the conditions of evidence possession, and it provides the wrong sort of response to the conflict between the Fallibility Norm and the Evidentialist Dictum (§4). The second defense attempts to replace the conception of evidential support the problem relies on with a suitable alternative that clearly accommodates fallibility considerations. I discuss the challenges this defense faces in §5. I conclude by suggesting that we should reconceive the Evidentialist Dictum's role in epistemic rationality and by sketching an account that reconciles it, once its role is reconceived, with the Fallibility Norm (§6).

One upshot of the paper is that a common and natural stance on the epistemic significance of fallibility can no longer be taken for granted. More often than not, when I discuss these issues with philosophers who haven't given much thought to the Fallibility Norm, these philosophers tend to express a deflationary, if not dismissive, attitude: "When I receive evidence of epistemic error, isn't that just more evidence to be balanced against all the rest?"[6] Although this was a compelling, pretheoretical stance, if this paper shows anything, it shows that things are much more complicated. Whether the Evidentialist Dictum can be reconciled with the Fallibility Norm depends on subtle and complex issues surrounding evidence and evidential support.

2. THE PROBLEM

We make all sorts of errors: we misremember, we think we see something we didn't, we make fallacious inferences, we misapply good rules of inference, etc. These are *epistemic* errors because they lead to inaccurate and/or irrational beliefs. Throughout, my reference to our fallibility signifies our proclivity toward epistemic error.

To begin to see the problem our fallibility gives rise to, let's start with an example:

> **Hypoxia:** After doing some calculations while piloting an airplane, Andy becomes highly confident that she has enough fuel to fly to Hawaii on the basis of the following evidence, *E*:

- A full tank contains 20,000 miles worth of fuel.
- The tank is ¾ full.
- Hawaii was 16,000 miles from Andy's point of departure.
- Andy has flown 5,000 miles toward Hawaii.

    Then air traffic control warns her that, as a result of a drop in her oxygen supply, any reasoning she's done in the last five minutes has probably been distorted without her knowing it.[7]

In this example, Andy receives evidence that she has likely made an error in her calculations. Call evidence of epistemic error *higher-order evidence*.[8] The widespread intuition that Andy should reduce her confidence is typically encoded in the judgment that maintaining her confidence after she receives the higher-order evidence would be epistemically irrational.[9] I will suppose that this judgment is correct, and that it represents the fact that Andy, like the rest of us, must take her fallibility into account when managing her beliefs.

To make the problem more precise, I will replace talk of the Evidentialist Dictum with talk of *Evidentialism*, the view that believing a proposition is epistemically rational for someone just in case her evidence supports that proposition. On my usage, Evidentialism implies that rational belief supervenes on evidential support, rather than evidence alone: any two people for whom the same propositions are supported (to the same degree) by their respective evidence are justified in believing exactly the same things.[10] Evidentialism is a broad view that includes both objective and subjective Bayesianisms, for instance.[11]

One additional claim will be used to generate the problem, *Entailment-Support*: If a body of evidence entails some empirical (contingent *a posteriori*) proposition, that body of evidence supports that proposition.[12] This claim is controversial. Some will view it as downright implausible, while others will think it is the *only* objective evidential support principle.[13] I'm not sure it's the only objective evidential support principle, but I'm confident that it's not downright implausible. The principle, for instance, does *not* say that if your body of evidence entails *any* proposition (e.g., necessary truths), your body of evidence supports it. Moreover, it should be endorsed by anyone who accepts a probabilistic account of evidential support, which I consider a plausible conception of evidential

---

7   This case, originally due to Elga (ms), is widely discussed in the literature. I've added details about Andy's evidence.

8   Although this term has different uses, I'm restricting my use to evidence of epistemic error.

9   On similar cases, see Christensen (2010, 2013), Elga (2013, ms), Schechter (2013), and Schoenfield (2015).

10 This leaves open whether permissivism is true because what one's evidence supports may depend on more than one's evidence. Papers on permissivism include White (2005, 2010), Feldman (2007), Christensen (2007), Titelbaum (2010), Meacham (2014), and Schoenfield (2014).

11 See Ballantyne and Coffman (2011) on different evidentialisms. Two additional notes. First, as stated, Evidentialism might be accepted by self-proclaimed Evidentialists like Conee and Feldman (2004). However, anyone sympathetic with their (2008) developments of Evidentialism may not find the problem described in this section troubling. In §5, I make remarks that must be addressed by someone sympathetic to Conee–Feldman-style Evidentialism. Second, one might question whether subjective Bayesianism, or any view on which epistemic rationality supervenes on anything besides evidence, counts as a version of "Evidentialism." I include subjective Bayesianism to show how far-ranging the problem I'm going to discuss is. Whether or not subjective Bayesianism is appropriately thought of as a version of Evidentialism, it faces the same issues, as I explain in fn. 15.

12 This principle should be seen as a corollary of Evidentialism, since Evidentialists owe us an account of evidential support, and this principle provides a sufficient condition.

13 Titelbaum (2010) explicitly takes the latter stance.

support. In any case, in §5, I'll discuss challenges Evidentialists face if they reject this principle. Until then, I'll suppose it's true.

The problem is that Evidentialism cannot accommodate the judgment that Andy's belief is irrational once she receives the higher-order evidence. This is because, by Entailment-Support, Andy's evidence supports this belief, since her evidence entails it. More carefully: Andy is in an *Entailment Case of higher-order defeat* – an *Entailment Case* for short. In an Entailment Case, at t_1, the subject's evidence entails some proposition, P, and she believes P on the basis of this evidence. Then, at later t_2, she acquires higher-order evidence that renders her belief in P irrational.[14] By Evidentialism and Entailment-Support, Andy's belief is rational at t_1, and by Evidentialism and the judgment that Andy's belief is irrational at t_2, Andy's belief is not supported by her evidence at t_2. If it's not supported by her evidence at t_2, then, by Entailment-Support, it's not entailed by her evidence at t_2. But entailment is monotonic: adding evidence to a body of evidence without removing any evidence leaves the entailment intact. Hence, her belief is entailed by her evidence at t_2, hence it is supported, hence it is rational. Thus, the judgment that Andy's belief is irrational when she acquires the higher-order evidence appears to be inconsistent with Evidentialism and Entailment-Support. Since the latter principle spells out when a body of evidence supports a proposition, this appears to be an instance of the conflict: one cannot always simultaneously respect one's evidence and respect one's fallibility. Without a solution to this problem, Evidentialism will be unable to accommodate our judgments about how evidence of fallibility bears on rational belief.[15]

3. EVIDENCE-LOSS

Abandoning Evidentialism at this point would be hasty, for nothing has been said about whether Andy *loses* evidence between t_1 and t_2. Some philosophers discussing related problems do address this step. For instance, Kelly (2013: 46) writes:

> If one's original evidence entails that p, then it seems like one's total evidence will *always* support the belief that p, no matter how much misleading testimonial evidence one subsequently acquires, **so long as the original evidence remains part of the total set.** [Bold emphasis mine.][16]

---

14 This is a case of *higher-order defeat* because receiving higher-order evidence that a belief likely resulted from a misevaluation of the evidence renders that belief irrational.

15 This problem will arise for anyone who endorses *Entailment-Rationality*: If your evidence entails some (contingent *a posteriori*) proposition, believing that proposition is (all-things-considered) epistemically rational for you. In conjunction with Entailment-Support, Evidentialism implies Entailment-Rationality. But one needn't endorse Evidentialism to endorse Entailment-Rationality. Even if we shouldn't think of subjective Bayesianism as a version of Evidentialism, subjective Bayesians will still face this problem because they're committed to (something like) Entailment-Rationality. Accordingly, the solutions I float in §3 and §5 should be of interest to them. Henceforth, I'll simply assume that subjective Bayesianism is a version of Evidentialism.

16 Christensen (2010: 195) briefly addresses this step, too. Although Kelly is addressing peer disagreement, whereas I'm focusing on fallibility more generally, he only confirms that fallibility considerations create problems for Evidentialism when he says (Kelly 2013: 47): "Of course, it is difficult to make sense of the idea that someone who possesses entailing evidence should invest less than maximal credence in the entailed proposition; indeed, orthodox theories of evidential probability would seem to rule this out (at least in cases in which the person is certain of the entailing evidence itself)."

What Kelly says here is true. But he doesn't defend the claim that in the relevant cases the subject does not lose evidence. Yet, the considerations adduced in the last section suggest an argument from Evidentialism to the conclusion that Andy *does* lose evidence. Call it the Evidence-Loss Argument.

1. Evidentialism: Believing P is rational for S at t just in case S's evidence supports P at t.
2. Entailment-Support: If S's evidence entails (empirical) P at t, then S's evidence supports P at t.
3. Entailment Case: At t_1, S's evidence entails P, and S believes P on the basis of this evidence. At t_2, S acquires higher-order evidence that renders S's belief irrational.
4. *Entailment-Loss principle*: S's evidence entails P at t_1 and not at t_2 only if S loses evidence between t_1 and t_2.
5. S's evidence doesn't support P at t_2. (By (1) and (3))
6. S's evidence doesn't entail P at t_2. (By (2) and (5))
7. Therefore, S loses evidence between t_1 and t_2. (By (3), (4), and (6))

The Entailment-Loss principle is the only premise not discussed so far. It follows from the monotonicity of entailment and a plausible understanding of evidence. Entailment relations between sets of propositions cannot change, but which set of propositions constitutes one's evidence *can* change. Thus, this argument suggests the *Evidence-Loss Account* of higher-order defeat, that acquiring higher-order evidence that one has improperly assessed or is not in a position to accurately assess one's evidence sometimes causes one to lose some of that evidence.[17]

In addition to the Evidence-Loss Argument, there are two reasons to believe the Evidence-Loss Account. First, it has a precedent in the literature. Williamson (2000: 219) describes an example of evidence-loss occasioned by acquiring evidence that one misassessed one's original evidence:

> I see one red and one black ball put into an otherwise empty bag, and am asked the probability that on the first ten thousand draws with replacement a red ball is drawn each time. I reply '$1/2^{10,000}$'. Part of my evidence is the proposition *e* that a black ball was put into the bag; my calculation relies on it. Now suppose that on the first ten thousand draws a red ball *is* drawn each time, a contingency which my evidence does not rule out in advance, since its evidential probability is non-zero. But when I have seen it happen, I will rationally come to doubt *e*.

Williamson concludes, "Thus *e* will no longer form part of my evidence." In this case, something very improbable occurs, given the subject's assessment of his evidence. So, upon making the observation, he loses evidence as a result of becoming reasonably worried that he misassessed this evidence. Thus, acquiring evidence of misassessment can result in evidence-loss.

---

17   To avoid confusion, we need to distinguish *evidence possession* from *support strength*. It might be natural to describe someone who has weaker evidence than someone else as having less evidence. But having weaker evidence is compatible with possessing more evidence. The sense of "having evidence" relevant to the Evidence-Loss Account has implications regarding support strength, but it is the stronger, literal notion of possessing evidence that is at issue.

Second, the Evidence-Loss Account explains the illegitimacy of dogmatically dismissing higher-order evidence in a way that parallels the standard account of the illegitimacy of the reasoning involved in the Dogmatism Paradox for knowledge,[18] and its two main competitors on this front are incompatible with Evidentialism. Suppose Andy were to reason as follows:

> If I were suffering from hypoxia, I probably would have made the wrong inference. But I made the right inference: I concluded that I have enough fuel to fly to Hawaii and that is just what my evidence supports – indeed, entails! So, I'm probably not suffering from hypoxia.

Several philosophers have cautioned against such dogmatic responses to higher-order evidence.[19] Since the hypoxia warning calls into question Andy's evidence assessment, appealing to that assessment to conclude she's not suffering from hypoxia seems illegitimate. A similar form of dogmatic reasoning in response to misleading evidence against knowledge has been widely discussed. Consider Harman's (1973: 148) statement of this reasoning:

> If I know that h is true, I know that any evidence against h is evidence against something that is true; so I know that such evidence is misleading. But I should disregard evidence that I know is misleading. So, once I know that h is true, I am in a position to disregard any future evidence that seems to tell against h.

What explains why these sorts of reasoning are illegitimate?

In the knowledge case, the standard criticism is that the reasoning incorrectly presupposes that once one has knowledge, one may always thereafter use it in one's reasoning. But acquiring misleading evidence can cause one to lose the knowledge one might have used to dismiss that evidence.[20] Evidentialists can claim something similar about Andy's reasoning. It's illegitimate because although her evidence previously entailed that she had enough fuel, this doesn't warrant her in disregarding further evidence. Getting that further evidence may change what her evidence entails by causing her to lose evidence. This is just what the Evidence-Loss Account implies.

Evidentialists cannot accept two alternative explanations found in the literature. First, there is Christensen's (2010, 2011) *Bracketing Explanation*, on which accounting for her higher-order evidence requires Andy to bracket her original reasons. This is a "virtual" rather than "genuine" evidence-loss account: Andy doesn't actually lose evidence, but what she is justified in believing after she acquires the higher-order evidence depends only on the evidence remaining after she brackets her original evidence. While discussing disagreement, Kelly (2013: 45) discusses the parallel between dogmatic reasoning involving higher-order evidence and the Dogmatism Paradox, and proposes what I'll call the *Non-Monotonic Explanation* of why this reasoning is illegitimate:

> [A]fter I add the fact that you believe as you do to my stock of evidence, it will no longer be reasonable for me to believe that p, given what is my total evidence. And if it's no longer reasonable for me to believe that p, then I lack any rational basis for inferring that your sincere testimony is misleading evidence.

---

18 See Harman (1973) and Kripke (1971/2011) on the Dogmatism Paradox.
19 For seminal discussions, see Christensen (2007, 2011), Elga (2007), Kornblith (2010).
20 Harman (1973: 148–9).

The idea is that the support enjoyed by a subject's belief before she was aware of the disagreement is lost, in the standard way, once she learns of the disagreement. As a result, subsequently appealing to that belief to dismiss her peer's belief is irrational.

All three explanations imply that the dogmatic reasoning is unreasonable because some of the premises of the dogmatic reasoning are unjustified. However, Evidentialists can accept neither the Bracketing Explanation nor the Non-Monotonic Explanation. The problem cases are Entailment Cases: the original evidence entails the relevant beliefs, hence the Evidentialist must say that the original evidence supports the beliefs. No matter how much evidence one brackets, and no matter how much additional evidence gets added to one's body of evidence, if one does not *lose* any evidence, then according to Evidentialism the premises of the problematic reasoning remain justified. Indeed, Christensen argues that Andy's evidence *does* still support her beliefs because, on his non-Evidentialist account, higher-order evidence defeats while leaving evidential support relations intact.[21] Although Kelly wants to provide an Evidentialist account, he admits that Entailment Cases give rise to a genuine puzzle.[22] The Evidence-Loss Account allows Evidentialists to solve that puzzle.

4. AGAINST THE EVIDENCE-LOSS ACCOUNT

The Evidence-Loss Account should strike Evidentialists and non-Evidentialists alike as a serious hypothesis. In spite of its attractions, however, it is open to at least three objections. The first shouldn't worry Evidentialists much. The second and third strike me as strong objections.

First, the Evidence-Loss Account says when Andy acquires the higher-order evidence she thereby loses evidence. But which evidence does she lose? No principled answer seems available. Given the description of the case, there are four main options, but picking one rather than the other is arbitrary.

This objection should be addressed if the Evidence-Loss Account is defended further. However, lacking a precise, general answer is not a major embarrassment, since the main alternative faces a similar question, and also lacks a precise, general answer. This alternative is Christensen's Bracketing Account, according to which, accommodating higher-order evidence requires bracketing some of one's original evidence.[23] But which evidence, or how much of it, must Andy bracket?[24] It's unlikely that any principled answer will be forthcoming. Thus, the Evidence-Loss Account can adopt a companions-in-guilt strategy. It's not off the hook. But this is a challenge rather than a strong objection.

According to the second objection, the Evidence-Loss Account requires an implausible "purpose-relative" conception of evidence possession. Consider a case of peer disagreement. Suppose that while out to dinner, you and your friend discover that you disagree about your shares of the bill. Even if, in some sense, you thereby lose evidence, you can

---

21   Christensen (2010: 195, 197).

22   Kelly (2013: 46). Kelly attempts to mitigate the problem posed by this puzzle by arguing that it arises independently of higher-order evidence considerations. I think he's mistaken about this. See DiPaolo (ms).

23   This account underlies Christensen's Bracketing Explanation (§3).

24   Christensen (2010: 213; 2011).

still legitimately appeal to this evidence for *some* purposes. Suppose it's part of your original evidence that the bill, before tip, is $60. You can reasonably reject your friend's claim that your share is $400, since the bill itself is only $60.[25] But, to avoid condoning dogmatism, the Evidence-Loss Account prohibits appealing to this evidence to dismiss the more plausible position that, with tip, your shares are $35, rather than $36, say. Now suppose you're out to dinner with two people: one claims your shares are $400 and the other claims $35. Arguably, you can reasonably dismiss the $400 verdict, but not the $35 verdict. What explains this? The natural answer is that for the purposes of ruling out the ridiculous response you can reasonably appeal to your original evidence, but for the purposes of ruling out the sensible response you cannot.[26] The Evidence-Loss Account explains the former fact by positing evidential support, hence evidence possession, and it explains the latter fact by positing evidence-loss. Thus, if the proponent of the Evidence-Loss Account wants to capture all of these facts, she seems committed to a purpose-relative account of evidence possession: you possess the evidence about the bill for the purpose of dismissing the ridiculous response, but not for the purpose of answering the sensible one.

While the claim that evidence possession is purpose-relative is not obviously incoherent, it seems implausible. It's not as if considerations available to you for the purpose of answering one question vanish when you attend to another. But on the Evidence-Loss Account they do lose their status as *evidence* depending on your purposes or the questions you're addressing. This means if you were aware of the Evidence-Loss Account's truth, things would end up looking very strange from your perspective. You might think: "I previously possessed evidence E, which supported belief in P and in Q. Now that I have this higher-order evidence, I *still* possess E for the purpose of answering whether P, E still bears on the question whether Q, but I do *not* possess E for the purpose of answering whether Q." This doesn't seem right. The Evidence-Loss Account introduces relativity into evidence possession where there appears to be none.

Of course, proponents of the Evidence-Loss Account can insist that this mischaracterizes which evidence has been lost and any trace of purpose-relativity will vanish once the proper characterization is in place. However, this response invites the question: Why does the subject lose *that* evidence rather than *this* evidence? The more heavily the account relies on particular judgments about which evidence one loses in response to

---

25   Cf. Christensen (2007).

26   Although Christensen (2007: 200–1) seems to implicitly endorse appealing to the original evidence in a case like this, his theoretical commitment to Independence may prohibit this appeal. (However, Lord (2014) and Kelly (2013) argue that the verdicts Christensen wants might require appealing to the original evidence in the way I'm suggesting.) Things are more complicated in Christensen (2011) because although he doesn't endorse this appeal in his discussion of similar cases, he (2011: 18) admits that there are unresolved issues about the scope of Independence. In any case, Evidentialists who endorse Independence may find my argument objectionable. But, the argument from the Evidence-Loss Account to purpose-relativity doesn't depend on this example. It only requires (i) that there is *some* purpose for which the agent who's suffered higher-order defeat can use her original evidence, but (ii) that she could use it for that purpose only if she possesses it. For instance, after providing her with the hypoxia warning, air traffic control might ask Andy how much fuel she has. It seems perfectly legitimate for her to appeal to her evidence that the tank is ¾ full. (Fill in the details however you like to focus on another piece of her evidence.) If that's right, then in order to capture this point, the Evidence-Loss Account requires a purpose-relative conception of evidence possession.

higher-order evidence, the more urgent the need to answer the first objection begins to look. I think placing much theoretical weight here would be unwise.

The purpose-relativity of evidence possession might be defended by appealing to the idea that epistemic notions are 'interest-relative,' in the sense that, for instance, whether a subject knows or is justified in believing something depends partly on her interests or what is at stake for her.[27] Indeed, Stanley (2005: 182) suggests that *all* epistemic notions are interest-relative, including *evidence*. Inspired by Stanley's position – *Strong Interest-Relativism* – a proponent of the Evidence-Loss Account may double down and defend the purpose-relativity of evidence possession.

Certainly by some people's lights we needn't take the Evidence-Loss Account seriously if it relies on Strong Interest-Relativism. But even if Strong Interest-Relativism is defensible, the claim that evidence possession is purpose-relative is even more radical. First, the Evidence-Loss Account is primarily concerned with evidence possession, rather than support strength. Some of Stanley's remarks seem to suggest that what is relative to one's interests is support strength, not possession. But that won't justify the purpose-relativity of evidence possession. Stanley also sympathetically discusses Williamson's E = K account of evidence, according to which, one's evidence is just what one knows. If E = K and Strong Interest-Relativism are true, then whether a proposition is part of one's evidence will depend on one's interests. Yet, this doesn't guarantee purpose-relativity. The Evidence-Loss Account allows evidence possession to vary while interests are held fixed, since interests and purposes are distinct. There may be much or little at stake for you, but once we hold fixed your interests, Strong Interest-Relativism will have fixed what evidence you possess, whereas the Evidence-Loss Account will still leave this purpose-relative. So, in this respect the Evidence-Loss Account is more radical than Strong Interest-Relativism.

According to the third objection, the Evidence-Loss Account gives the wrong sort of response to Evidentialism's problem with accommodating higher-order defeat. On the Evidence-Loss Account, Andy's belief turns out to be irrational because she loses evidence. But a common concern that arises when we consider our fallibility is that we make "performance" errors while attempting to properly form beliefs. Hume (1978: 180), for instance, insists that although deductive rules of inference are certain and infallible, our *application* of them is not.[28] Likewise, when Andy gets the hypoxia warning, this is worrisome because it makes it likely that she made a faulty inference. And, it seems, this sort of worry leaves an agent's evidence untouched.

Compare Williamson's ball example (§3) with the Hypoxia example. In Williamson's case, the subject should abandon belief in his evidence because what he observes is incredibly unlikely if his evidence is true. Since incredibly unlikely things are probably false, his evidence is probably false, too. This is why claiming evidence-loss is plausible. But nothing like this occurs in Hypoxia. Andy takes her evidence to support a proposition, while having strong evidence that her take on the evidence is flawed. Unlike the subject in Williamson's case, what Andy observes is *not* incredibly unlikely if her evidence is true. Putting the comparison slightly differently: in contrast to the subject in Williamson's example who acquires reason to think he has misassessed what his evidence *is*, Andy only acquires reason to think she has misassessed what her evidence *supports*. Whereas

---

27 See Hawthorne (2005), Stanley (2005), and Fantl and McGrath (2009).

28 Schechter (2013: 448) makes a similar point.

continuing to accept what she took her evidence to support when she gets the hypoxia warning is unreasonable, relying on her judgments about what evidence she possesses seems unproblematic. Thus, while it's true that Williamson's example sets a precedent for the Evidence-Loss Account, this account doesn't properly explain all problematic Entailment Cases.

Since the Evidence-Loss Account arguably leads to an implausible, purpose-relative conception of evidence possession, and it provides the wrong sort of response to the problem raised by Entailment Cases insofar as that problem is about performance errors, I think we have good reason to reject the account as implausible and/or ultimately unmotivated.

5. EVIDENTIAL SUPPORT

The most natural move for Evidentialists to make at this point is to reject Entailment-Support.[29] By itself, rejecting this principle wouldn't show that Evidentialism *can* capture higher-order defeat, but it would allow Evidentialists to avoid the argument based in Entailment Cases that Evidentialism *can't* capture higher-order defeat.

For several reasons, I think we shouldn't find this satisfying. First, Evidentialists who reject Entailment-Support must abandon a plausible conception of evidential support because something like Entailment-Support should be endorsed by anyone who thinks of evidential support probabilistically. Given certain Evidentialists' (e.g., Bayesians') other theoretical commitments, rejecting Entailment-Support and/or a probabilistic conception of evidential support won't be a plausible option. More than this, though, the probabilistic conception strikes me as plausible. One of the main reasons it's plausible to think, as Evidentialists do, that rational belief is based in evidential support is that typically the more epistemically rational a belief is, the more likely it is to be true, and, likewise, typically the greater evidential support one has for a belief, the more likely it is to be true.[30] When we ask about the strength of the evidence we are asking how probable a hypothesis is given that evidence. This suggests that evidential support should be understood probabilistically.[31]

Evidentialists might reject the probabilistic conception in general, and Entailment-Support in particular, by appealing to human limitations. If Evidentialism is wedded to Entailment-Support, Evidentialism implies that we are rational in believing infinitely many propositions.[32] But some think of rationality *deontologically*, that is, in

---

29 One might grant Entailment-Support, but argue as follows: Evidentialists needn't worry about Hypoxia because although entailing evidence supports propositions, the specification of Andy's evidence is too broad. As specified, Andy's evidence includes propositions like *that the tank is ¾ full*, but it ought to be specified more narrowly, i.e., internally – e.g., *that the fuel gauge seems to read ¾ full*. However, this concern about evidence's *extent* doesn't help Evidentialists avoid the problem because the problem doesn't depend on a particular account of evidence's extent. Even if Andy's evidence is specified more narrowly, insofar as evidence is propositional, that evidence will entail some proposition, and she can acquire evidence that she bungled her reasoning about *that* proposition. Evidentialists would face the same problem.

30 See Goldman (1979), Williamson (2000, 202), and Sliwa and Horowitz (2015).

31 There are certainly other reasons to think of evidential support probabilistically. However, I will devote most of this section to addressing Evidentialist responses that reject this approach.

32 This is merely illustrative. My response doesn't depend on which cognitive limitations are appealed to.

terms of epistemic obligation, oughts, or requirements.[33] Moreover, many endorse an ought-implies-can principle in epistemology.[34] Since humans are finite, we cannot believe infinitely many propositions. So, the argument might go, we are not obligated to do this, so Entailment-Support and probabilistic conceptions of evidential support aren't true.[35]

There's plenty to question in this argument. One response, of course, is to insist that since Entailment-Support is true, Evidentialism is false. Setting that aside, though, it's not clear that Evidentialists should rely on this argument. Not only does nothing require the combination of views the argument relies on, it may also be true that this combination should be avoided. Alston (1988) argues against deontological conceptions of justification by appealing to ought-implies-can, whereas Feldman (2000) – himself, an Evidentialist – defends a deontological conception by arguing against ought-implies-can. One plausible conclusion to draw from these arguments, independently of your stance on Evidentialism, is that ought-implies-can shouldn't be paired with the deontological approach.[36] If that's right, then the argument relies on a conjunction that should be eschewed, and Evidentialists cannot reasonably defend their rejection of Entailment-Support in this way.

The second reason we shouldn't be satisfied with a simple rejection of Entailment-Support is that Evidentialists may be unable to capture higher-order defeat even if Entailment-Support is false. Christensen (2010: 197) argues that even in non-Entailment Cases higher-order evidence defeats while leaving evidential support intact. Suppose my evidence strongly supports, but does not entail, some hypothesis, and I adopt belief in that hypothesis on the basis of an inference to the best explanation, only to learn from a trustworthy source that I was slipped a reason-distorting drug before I made the inference. Christensen argues that I cannot rationally maintain my belief even though my evidence still supports it. Why does my evidence still support the belief? Because the evidence is not in question (as in Williamson's case). And the explanatory connection between that evidence and my belief remain incredibly strong, since these connections don't depend on any claims about me or my current reasoning abilities.

Christensen's argument depends on the conjunction of (1) H is the best explanation of E_1 and (2) H is the best explanation of $E_1 \& E_2$. He can stipulate (1), but whether (2) is true depends on the case. Since some instances of this conjunction are false, Evidentialists might reasonably claim that Christensen's case – where E_2 is the evidence about the reason-distorting drug – is one such instance. Evidentialists, like Conee and Feldman (2008), who think evidential support relations are explanatory relations may insist that the higher-order evidence shifts evidential support by changing the best explanation.

For what it's worth, it doesn't seem that this will always be true. Imagine Darwin adopting belief in the theory of evolution after surveying all of his data, only to be told that he'd just ingested a reason-distorting drug right before he had reasoned to this conclusion. Is it really plausible to say that the theory of evolution is no longer the best explanation of his data? If we set the case up right, it will certainly seem irrational for him to maintain his belief. But it's not obvious to me that evolution will no longer be the best explanation.

---

33 See Feldman (2000).
34 For example, Lycan (1975, 1985), Alston (1988), Dretske (2000), and Greco (2012).
35 See Goldman (1978) for related arguments.
36 Elsewhere, I consider the relation between cognitive limitations and ought-implies-can. The details matter here.

In any case, without principles specifying the circumstances in which a hypothesis is the best explanation of some body of evidence, this disagreement will boil down to a clash of judgments about particular cases. Christensen thinks, at least in a range of cases, acquiring higher-order evidence doesn't change the fact that the relevant hypothesis is the best explanation. Conee and Feldman, and other Evidentialists, might disagree.[37] It's true that Evidentialists can avoid the problem by altogether abstaining from providing principles of evidential support, so that no determinate, principled judgments can be made about whether, on their view, acquiring some evidence leaves evidential support intact. But Evidentialism would be much more secure – i.e., we could be more confident that Evidentialism *can* capture higher-order defeat – if Evidentialists addressed the problem by presenting an account of evidential support that rendered the right verdicts without running into additional problems.

Evidentialists might look to their opponents for help on this front. Presumably, non-Evidentialists think there's some principled way to explain why Andy should revise her belief. Depending on the explanation, Evidentialists may be able to convert this principle about rational belief into a plausible principle of evidential support.[38] For example, suppose the following principle were advanced by non-Evidentialists to explain why Andy should revise her belief:

> **Levels**: It's irrational to believe P when you think that your evidence doesn't support P.[39]

The Evidentialist might convert Levels into a principle of evidential support:

> **Levels-Support**: Your evidence doesn't support P when you think your evidence doesn't support P.

More generally, it might seem that most stories the non-Evidentialist would tell about higher-order defeat could just as plausibly be told by Evidentialists in terms of evidential support.[40] We can call this the *Conversion Strategy*.[41]

The Conversion Strategy faces at least two challenges. First, if applied indiscriminately, the Conversion Strategy runs the risk of trivializing Evidentialism, a substantive thesis. To see this concern, first note that some argue evidential support and coherence – more generally, reasons-based normativity and wide-scope rationality-based normativity – are distinct. Accordingly, some argue that coherence, not evidential support, is the source of

---

37   Kelly's (2010) Total Evidence View will also fall in this category. He captures higher-order defeat only by refraining from providing principles of evidential support.

38   After Schoenfield (2015) discusses various attempts to account for higher-order defeat, she concludes that the relevant principles should *not* be understood as principles about the relationship between bodies of evidence and belief states, but should instead be understood as principles of "reasoning." Her arguments may cast additional doubt on the proposal I'm considering.

39   Cf. Schechter (2013: 443): "I should be much less confident in the result of my directional calculation if I think it likely that I have hypoxia." Also, cf. Christensen (2013: 91): A "level-connection" ideal precludes one "from being highly confident of P while simultaneously believing that that high degree of confidence is much higher than that supported by one's evidence."

40   Thanks to an anonymous Associate Editor at *Episteme* for suggesting that I address this response.

41   The relation the Conversion Strategy bears to Evidentialism is analogous to the relation the "Consequentializing" strategy bears to Consequentialism. See Portmore (2009) for an overview of Consequentializing.

higher-order evidence's normative force.[42] Now consider Levels and Levels-Support. Levels appears to be a coherence principle, whereas Levels-Support is, by design, an evidential support principle. By converting Levels into Levels-Support, Evidentialists blur the distinction between these two possible sources of normativity. Maybe Evidentialists are right to do this in this case. But in order to avoid trivializing Evidentialism, Evidentialists had better recognize *some* alleged source of normativity that cannot be converted into evidential support.

Even if the Conversion Strategy can be applied without trivializing Evidentialism, it still faces a second challenge. Recall from the introduction that several philosophers object to conciliationism on the grounds that it recommends ignoring first-order evidence. Nothing is special about conciliationism here. To respect our fallibility, prominent non-Evidentialist accounts require us to downplay, one way or another, the significance of our first-order evidence. Non-Evidentialists can live with this because their focus is directly on rational belief, rather than evidential support. But if Evidentialists convert non-Evidentialist accounts into principles of evidential support, Evidentialists will thereby transfer this feature (or bug) into their account of evidential support. This might make them susceptible to the sort of objection often leveled against their opponents. For instance, if part of the complaint against conciliationism is that it fails to give the first-order evidence its proper weight, and Evidentialists convert conciliationism's principles of rational belief into principles of evidential support, then Evidentialists will have thereby adopted principles of evidential support that conflict with common judgments about first-order evidence's proper role in evidential support. This would be bad. One would have thought if any view could give the first-order evidence its due, Evidentialism could. If that's right and Evidentialism must ultimately abandon certain claims about first-order evidence's role in determining total evidential support, then a standard argument against non-Evidentialist positions, and in favor of Evidentialism will be undermined.

Thus, the challenge for Evidentialists who want to pursue the Conversion Strategy is to find principles of rational belief (1) that accommodate higher-order defeat, (2) that can be converted into plausible principles of evidential support, and (3) that give the first-order evidence its due. Although I have no knockdown argument that this challenge cannot be met, I'll consider two implementations of the Conversion Strategy to illustrate the difficulty of the challenge.

Begin with Levels-Support. Recall, it says an agent's evidence doesn't support a proposition if she thinks it doesn't. It's based on Levels, which says rationality forbids you from believing a proposition when you believe your evidence doesn't support that proposition. According to an Evidentialism that includes Levels-Support, Andy should revise her belief because having acknowledged that her evidence probably doesn't support her belief that she has enough fuel, her evidence therefore no longer supports this conclusion. That is, her evidence no longer supports this conclusion *because* she thinks it doesn't. But this makes evidential support far too subjective. The first-order evidence seems to drop out of the picture completely. Moreover, it implies that anyone can rationally resist counterevidence simply by being dogmatic or pigheaded about what their evidence supports. Of course, Evidentialists might question whether Levels-Support is any worse off than Levels here. But I think it is. Levels might be paired with a more objective account of evidential support, in which case even if the pigheaded person's beliefs satisfy Levels, his beliefs might be

---

rationally criticizable for totally failing to be sensitive to evidence. If Evidentialism and Levels-Support are both true, the pigheaded person's beliefs will be rationally uncriticizable.

A more sophisticated approach faces a related problem. Sliwa and Horowitz (2015) argue that accounts of higher-order evidence typically fail to respect *all* of the evidence, by failing to give proper weight either to first-order evidence or to higher-order evidence.[43] They explicitly set out to defend a view that avoids this failure. Here's their view:

> **Evidential Calibration:** When one's evidence favors P over ∼P, one's credence in P should equal the expected reliability of one's educated guess that P.[44]

Applying Evidential Calibration requires three steps. First, determine which proposition your first-order evidence favors, where evidence *favors* P over ∼P just in case the degree of support it provides to P is higher than the degree of support it provides to ∼P. Second, determine your expected reliability, r, in making an educated guess about that proposition. Finally, assign r to the favored proposition. If your evidence favors P and your expected reliability is .73, then according to Evidential Calibration, if you have a credence of .73 in P, your belief is rational and you have accounted for all of your evidence.

First-order evidence doesn't completely drop out of Evidential Calibration's picture, but if we interpret it as a principle of evidential support, we do get some odd results. Suppose at t_0, your rational credences are 50/50 in P and ∼P, at t_1 you get evidence that slightly confirms P, and then at t_2 you get slightly disconfirming evidence. Suppose your expected reliability about P remains .99 across these times, while the degree of support P receives from your first-order evidence, given the slightly confirming and disconfirming evidence, is .51 and .49 at t_1 and t_2, respectively. Understood as a principle of evidential support, Evidential Calibration implies that the total degree of your evidential support for P goes from .5 to .99 to .01 as you receive what appears to be nearly negligible first-order evidence. Consider a second case. Suppose your expected reliability remains .89 from t_1 to t_3, but you receive first-order evidence that shifts the first-order degree of support for P from .51 up to .99 back down to .51. According to Evidential Calibration, your total degree of evidential support remains at .89 throughout this interval. On this view, then, slight shifts in first-order evidential support lead to huge swings in total evidential support whereas large shifts in first-order evidential support lead to no changes in total evidential support. This doesn't look like a view that takes first-order evidence seriously.

As Sliwa and Horowitz (2015: 2854) explain, Evidential Calibration "takes the *valence* of our first-order evidence into account – whether the evidence favors P or ∼P – [but] it ignores the *strength* of that evidence." Thus, Evidentialists who apply the Conversion Strategy to Evidential Calibration will be forced to say that when agents have higher-order evidence, the only role first-order evidence plays in determining evidential support is setting valence. Given the two cases just described, this role seems at once too significant and not significant enough. And it's worth emphasizing that Sliwa and Horowitz crafted

---

43   Indeed, Sliwa and Horowitz (2015: 2840) level the "fails to give the first-order evidence its due" objection that I've been discussing against a view developed in White (2009) that is similar to Levels.

44   I'll interpret this view slightly differently from how Sliwa and Horowitz do to avoid some orthogonal complexities. See Sliwa and Horowitz (2015: 2843–5) for the details.

Evidential Calibration with the explicit intention of giving first-order evidence its due. Yet, when it's converted into a principle of evidential support, it leads to questionable results, at best.

Let's sum up. If Evidentialists want to capture higher-order defeat, they'll need to reject Entailment-Support, and a probabilistic conception of evidential support along with it. But to ensure that they can accommodate higher-order defeat, rather than simply avoid the argument that they can't, they need to do more than deny Entailment-Support. Refusing to provide principles of evidential support leaves Evidentialism on shaky grounds. The Conversion Strategy may be a promising way forward for Evidentialism. But two implementations of this strategy have implausible implications about evidential support because they both fail to give first-order evidence its due. While diminishing the impact of first-order evidence in the ways these theories do may be appropriate for a theory of rational belief, it seems inappropriate for a theory of evidential support. Going forward, the task for Evidentialists is to defend one of these implementations or offer a distinct theory of evidential support that provides both the right verdicts about evidential support and about higher-order defeat. Either way, Evidentialists have their work cut out for them.

6. CONCLUSION

We began with an apparent tension between the Evidentialist Dictum and the Fallibility Norm. I operationalized these norms by focusing on Evidentialism and higher-order defeat, respectively. If what I've argued is correct, Evidentialism has trouble accommodating higher-order defeat. The tension looks genuine. However, I think our theorizing should make room for both the Evidentialist Dictum and the Fallibility Norm.[45]

The Evidentialist Dictum and the Fallibility Norm are different kinds of norms. To illustrate, compare the following pair of norms from ethics: You ought not murder; if you murder, you ought to murder gently. We can call the first norm a *primary* norm and the second norm either a *secondary* norm or a *contrary-to-duty* norm because it tells you how you ought to act when you fail to act as you ought to (according to the relevant primary norm).[46] Contrary-to-duty norms are *norms of compensation*: they tell us how to compensate for a certain sort of imperfection, namely the imperfection of failing to act as we ought. There are other norms of compensation, as well. Unlike contrary-to-duty norms, some of these other norms of compensation are sensitive not to actual compliance failure but to the *risk* of compliance failure. Even if a doctor ought to administer drug A, if there's a chance that she'll fail to comply with this norm, then it might be that she ought to administer a different drug B. For instance, suppose drug A would be best, drug B would be second best but still acceptable, and drug C would be lethal. And suppose all of her pills of drug A have been mixed with her indistinguishable pills of drug C, while her pills of drug B are separate and distinguishable. In this case, she *could* comply with the primary norm (by successfully grabbing a pill of drug A at random), but the risk of failing to comply is high and the consequences of failing to comply while trying to comply are grave. Given this risk, she ought to fail to comply with the primary norm in a different way, namely by administering drug B. Likewise, I think the Fallibility Norm is a norm of compensation sensitive to the risk of failing to comply

---

45   See DiPaolo (ms) for more details.
46   See Chisholm (1963) on contrary-to-duty norms.

with the Evidentialist Dictum, a primary norm. Our fallibility puts us at risk of failing to comply with the Evidentialist Dictum, and to compensate for that fact we need to be sensitive to higher-order evidence, which in turn can require us to fail to comply with the Evidentialist Dictum. Just as we can maintain the compatibility of primary norms and norms of compensation in spite of the different recommendations they issue, we can likewise maintain the compatibility of the Evidentialist Dictum and the Fallibility Norm.[47]

REFERENCES

Alston, W. 1988. 'The Deontological Conception of Justification.' *Philosophical Perspectives*, 2: 257–99.
Ballantyne, N. and Coffman, E.J. 2011. 'Uniqueness, Evidence and Rationality.' *Philosophers' Imprint*, 11: 1–13.
Chisholm, R. 1963. 'Contrary-to-Duty Imperatives and Deontic Logic.' *Analysis*, 24: 33–6.
Christensen, D. 2007. 'The Epistemology of Disagreement: The Good News.' *Philosophical Review*, 116: 187–217.
—— 2010. 'Higher-Order Evidence.' *Philosophy and Phenomenological Research*, 81: 185–215.
—— 2011. 'Disagreement, Question-Begging, and Epistemic Self-Criticism.' *Philosopher's Imprint*, 11: 1–22.
—— 2013. 'Epistemic Modesty Defended.' In D. Christensen and J. Lackey (eds) *The Epistemology of Disagreement: New Essays*, pp. 77–97. New York, NY: Oxford University Press.
Conee, E. and Feldman, R. 2004. *Evidentialism*. New York, NY: Oxford University Press.
—— and —— 2008. Evidence. In Smith Q. (ed.) *Epistemology: New Essays*, pp. 83–104. New York, NY: Oxford University Press.
DiPaolo, J. 2016. 'Higher-Order Defeat is Object-Independent.' *Pacific Philosophical Quarterly*. doi: 10.1111/papq.12155.
—— ms. 'Fallibility and Normativity.'
Dretske, F. 2000. 'Entitlement: Epistemic Rights without Epistemic Duties?' *Philosophy and Phenomenological Research*, 60: 591–606.
Elga, A. 2007. 'Reflection and Disagreement.' *Noûs*, 41: 478–502.
—— 2010. 'How to Disagree about How to Disagree.' In R. Feldman and T. Warfield (eds), *Disagreement*, pp. 175–86. New York, NY: Oxford University Press.
—— 2013. 'The Puzzle of the Unmarked Clock.' *Philosophical Studies*, 164: 127–39.
—— ms. 'Lucky to be Rational.'
Fantl, J. and McGrath, M. 2009. *Knowledge in an Uncertain World*. New York, NY: Oxford University Press.
Feldman, R. 2000. 'The Ethics of Belief.' *Philosophy and Phenomenological Research*, 60: 667–95.
—— 2005. 'Respecting the Evidence.' *Philosophical Perspectives*, 19: 95–119.
—— 2007. 'Reasonable Religious Disagreement.' In L. Antony (ed.), *Philosophers without Gods*, pp. 194–214. New York, NY: Oxford University Press.
—— 2009. 'Evidentialism, Higher-Order Evidence, and Disagreement.' *Episteme*, 6: 294–312.
Field, H. 2000. 'Apriority as an Evaluative Notion.' In P. Boghossian and C. Peacocke (eds), *New Essays on the A Priori*, pp. 117–49. Oxford: Oxford University Press.
Goldman, A. 1978. 'Epistemics: The Regulative Theory of Cognition.' *Journal of Philosophy*, 75: 509–23.

—— 1979/1992. 'What is Justified Belief?' In G. Pappas (ed.), *Justification and Knowledge*. Reprinted in *Liaisons*, 1992, pp. 105–26. Cambridge, MA: MIT Press.

Greco, D. 2012. 'The Impossibility of Skepticism.' *Philosophical Review*, 121: 317–58.

Harman, G. 1973. *Thought*. Princeton, NJ: Princeton University Press.

Hawthorne, J. 2005. *Knowledge and Lotteries*. New York, NY: Oxford University Press.

Hume, D. 1978. *A Treatise of Human Nature*, 2nd edition. P. H. Nidditch (ed.). Oxford: Oxford University Press.

—— 1999. *An Enquiry Concerning Human Understanding*. T. L. Beauchamp (ed.). Oxford: Oxford University Press.

Kelly, T. 2005. 'The Epistemic Significance of Disagreement.' *Oxford Studies in Epistemology* vol. 1. Oxford: Oxford University Press.

—— 2010. 'Peer Disagreement and Higher Order Evidence.' In R. Feldman and T. Warfield (eds), *Disagreement*, pp. 111–74. New York, NY: Oxford University Press.

—— 2013. 'Disagreement and the Burdens of Judgment.' In D. Christensen and J. Lackey (eds), *The Epistemology of Disagreement: New Essays*, pp. 31–53. New York, NY: Oxford University Press.

Kornblith, H. 2010. 'Belief in the Face of Controversy.' In R. Feldman and T. Warfield (eds), *Disagreement*, pp. 29–52. New York, NY: Oxford University Press.

Kripke, S. 1971/2011. 'Two Paradoxes of Knowledge.' In *Philosophical Troubles: Collected Papers*, Vol. 1. New York, NY: Oxford University Press.

Lasonen-Aarnio, M. 2014. 'Higher-Order Evidence and The Limits of Defeat.' *Philosophy and Phenomenological Research*, 88: 314–45.

Lewis, D. 1971. 'Immodest Inductive Methods.' *Philosophy of Science*, 38: 54–63.

Lord, E. 2014. 'From Independence to Conciliationism: An Obituary.' *Australasian Journal of Philosophy*, 92: 365–77.

Lycan, W. 1975. 'Occam's Razor.' *Metaphilosophy*, 6: 223–37.

—— 1985. 'Epistemic Value.' *Synthese*, 64: 37–64.

Meacham, C.J.G. 2014. 'Impermissive Bayesianism.' *Erkenntnis*, 79: 1185–217.

Portmore, D. 2009. 'Consequentializing.' *Philosophy Compass*, 4: 329–47.

Schechter, J. 2013. 'Rational Self-Doubt and the Failure of Closure.' *Philosophical Studies*, 163: 429–52.

Schoenfield, M. 2014. 'Permission to Believe: Why Permissivism is True and What it Tells Us about Irrelevant Influences on Belief.' *Noûs*, 48: 193–218.

—— 2015. 'A Dilemma for Calibrationism.' *Philosophy and Phenomenological Research*, 91: 425–55.

Sliwa, P. and Horowitz, S. 2015. 'Respecting *All* the Evidence.' *Philosophical Studies*, 172: 2835–58.

Stanley, J. 2005. *Knowledge and Practical Interests*. New York, NY: Oxford University Press.

Titelbaum, M. 2010. 'Not Enough There There.' *Philosophical Perspectives*, 24: 477–528.

—— 2015. 'Rationality's Fixed Point (Or: In Defense of Right Reason).' *Oxford Studies in Epistemology*, vol. 5. Oxford: Oxford University Press.

Weatherson, B. 2013. Disagreements, Philosophical and Otherwise. In D. Christensen and J. Lackey (eds), *The Epistemology of Disagreement: New Essays*, pp. 54–73. New York, NY: Oxford University Press.

—— ms. 'Do Judgments Screen Evidence?'

White, R. 2005. 'Epistemic Permissiveness.' *Philosophical Perspectives*, 19: 445–59.

—— 2009. 'On Treating Oneself and Others as Thermometers.' *Episteme*, 6: 233–50.

—— 2010. 'Evidential Symmetry and Mushy Credence.' *Oxford Studies in Epistemology*, vol. 3. Oxford: Oxford University Press.

Williamson, T. 2000. *Knowledge and its Limits*. New York, NY: Oxford University Press.

Worsnip, A. 2015. 'The Conflict of Evidence and Coherence.' *Philosophy and Phenomenological Research*. doi: 10.1111/phpr.12246.

JOSHUA DIPAOLO is a Visiting Assistant Professor in the Philosophy Department at Kansas State University. His research interests include epistemology, ethical theory, and political philosophy. His work has appeared in *Philosophical Studies*, *Synthese*, *Pacific Philosophical Quarterly*, and the *Journal of Ethics and Social Philosophy*.