

The Epistemic Function of Higher-Order Evidence

Declan Smithies

Draft chapter for *Propositional and Doxastic Justification*,
Edited by Paul Silva and Luis Oliveira (Routledge, Forthcoming).

Imagine you're a solo pilot flying a small plane in Arizona.¹ You're wondering whether to take a scenic detour via the Grand Canyon en route to your final destination. You know how far you can fly on a full tank and how much fuel remains in the tank. You also know the distance from here to the Grand Canyon and from there to your final destination. But you need to do some mental arithmetic in order to calculate whether you have enough fuel to safely complete the journey. You perform the calculation correctly and deduce that you have enough fuel. On that basis, you decide to take the detour. So far, so good!

Moments later, however, you acquire worrying new evidence that you are suffering from *hypoxia* – an oxygen deficit that impairs cognitive functioning in ways that often remain undetectable to the victim. You know about the dangers of hypoxia: pilots have crashed and died as a result of bad decisions made under its influence. You know there is a serious risk of hypoxia when the altitude is high enough and the cabin pressure is low enough. Moreover, your control panel says you're now in the danger zone. As it happens, this evidence is misleading, since there is a malfunction in the barometer that measures cabin pressure. The truth is that you're at no risk of hypoxia, although there's no way you can know this.

Should you now reconsider your decision to take the scenic detour once you acquire this new evidence? Intuitively, you should. Consider the *steadfast pilot* who decides to stick with his original plan. This decision seems grossly irrational. After all, the pilot has no way of knowing that the new evidence of hypoxia is misleading. Moreover, this evidence makes it likely that he is cognitively impaired in ways that dispose him to make elementary mistakes in calculation. It seems reckless to ignore this possibility in dogmatically maintaining his original plan. We know there is no risk of hypoxia, of course, and so we can be assured that the pilot will arrive safely. But any pilot who routinely makes such decisions in the face of evidence of hypoxia is putting their own life in serious danger.

The irrationality of the decision reflects the irrationality of the belief on which it is based. In general, it is rationally permissible to believe that p only if it's rationally permissible to act on the premise that p . And yet it's no longer rationally permissible to *act* on the premise that you have enough fuel when you acquire the new evidence that you're hypoxic. Hence, it's no longer rationally permissible to *believe* that you have enough fuel. This means that you cannot *know* whether you have enough fuel, since it's rationally permissible to believe that p , and to act on

¹ This example is adapted from Elga (2008, 2013) and Christensen (2010).

the premise that p , whenever you know that p .² Previously, however, you knew that you had enough fuel by competently deducing this conclusion from known premises. Therefore, you must lose this knowledge when you acquire the new evidence that you are hypoxic.

This much is intuitively compelling but theoretically puzzling. Why should the evidence that you are hypoxic destroy your knowledge that you have enough fuel? It is extremely plausible *that* it does, but it is not easy to explain *how* it does. After all, the evidence that you are hypoxic doesn't bear directly on the question of whether you have enough fuel: it isn't evidence that you *don't* have enough fuel. But then how exactly does this new evidence undermine the rationality of believing that you have enough fuel and thereby destroy your knowledge?

This is one instance of a more general question about the epistemic function of higher-order evidence. In this context, *higher-order evidence* is defined as evidence about whether your beliefs are responsive to your evidence. Suppose you know that p on the basis of evidence e , but then you subsequently acquire misleading higher-order evidence h that your belief is not responsive to your evidence. How does this higher-order evidence h destroy your knowledge that p ? And why is it now rationally impermissible for you to retain your belief that p on the basis of your first-order evidence e ?

This chapter provides a critical overview of several influential proposals in the literature on higher-order evidence. I start by criticizing explanations that appeal to evidential defeat (§1), epistemic conflicts (§2), and unreasonable knowledge (§3). Next, I propose an alternative explanation that appeals to a combination of improper basing (§4) and non-ideal rationality (§5). I conclude by summarizing my reasons for preferring this explanation to the alternatives (§6).

1. Evidential Defeat

Richard Feldman (2005) argues that higher-order evidence functions as an *evidential defeater*: it undermines knowledge by defeating evidence. For example, the higher-order evidence that you are hypoxic undermines your knowledge that you have enough fuel by defeating your first-order evidence for this conclusion. You lose your knowledge because your total body of evidence no longer supports this conclusion.

The challenge is to explain why the higher-order evidence that you are hypoxic defeats your first-order evidence that you have enough fuel. As we've already noted, your higher-order evidence doesn't provide evidence *against* this conclusion, since hypoxia by itself makes it no more or less probable that you have enough fuel. But then how can it defeat your evidence for this conclusion?

John Pollock (1986) draws an influential distinction between two kinds of defeaters. First, and most obviously, there are *rebutting defeaters*, which reduce

² Hawthorne and Stanley (2008) endorse both directions of this knowledge-action principle – that is, it's rationally permissible to act on the premise that p if and only if you know that p – but the opposite direction is more controversial.

your overall degree of evidential support for a conclusion by giving you evidence against the conclusion. Crucially, though, not all defeaters work this way. There are also *undercutting* defeaters, which reduce your overall degree of evidential support for a conclusion without giving you evidence against your conclusion. As Pollock writes, “Such defeaters attack the connection between the reason and the conclusion rather than attacking the conclusion itself” (1986: 196).

Consider a textbook example: the fact that the wall looks red is evidence that that it is red, although this evidential connection can be undercut by further evidence that the wall is bathed in red light. This is no evidence *against* the conclusion that the wall is red, since the lighting by itself makes it no more or less probable that the wall is red. As I’ll explain, however, it reduces the degree to which your sensory evidence supports this conclusion.

Under normal circumstances, the most probable explanation of why the wall looks red is that it is red. Hence, the evidence that the wall looks red raises the probability that it is red. Given the further evidence that the wall is bathed in the red lighting, however, this is no longer the most probable explanation. In these abnormal lighting conditions, the wall looks red whatever color it is. Hence, the evidence that the wall looks red under red lighting doesn’t raise the probability that the wall is red. Putting these points together, the probability that the wall is red given that it looks red under red lighting is less than the probability that the wall is red given only that it looks red. In this way, the evidence about the lighting reduces the degree to which your sensory evidence supports this conclusion. This is presumably what Pollock intends when he says that undercutting defeaters “attack the connection” between the evidence and the conclusion, rather than attacking the conclusion itself.

Now let’s revisit the hypoxia example: how does the evidence that you are hypoxic defeat your evidence that you have enough fuel? It’s not a rebutting defeater, of course, but not all defeaters are rebutting defeaters. Is it perhaps an undercutting defeater that reduces the degree to which your evidence supports this conclusion? I’ll argue that it isn’t, since the evidence that you’re hypoxic makes no difference to your overall degree of evidential support for the conclusion that you have enough fuel. There is no sense in which your higher-order evidence defeats the evidential support provided by your first-order evidence.

The key point is that the conclusion that you have enough fuel is *entailed* by known premises about the distance of the journey and the amount of fuel in the tank. Entailment is *monotonic*: we cannot undermine an entailment from premises to conclusion by adding new premises. In particular, we cannot undermine the entailment from premises to conclusion by adding the premise that you botched the deduction as a result of hypoxia. The expanded set of premises continues to entail the conclusion that you have enough fuel.

Moreover, entailment is the strongest kind of evidential support relation. Arguments come in varying degrees of strength: the stronger an argument, the higher the probability that its conclusion is true given that its premises are true. A deductively valid argument is the strongest kind of argument, since it is not merely

improbable but *impossible* that the conclusion is false when the premises are true. This is the limiting case in which the probability of the conclusion given the premises is 1. Since the premises of a deductively valid argument entail its conclusion, the probability of the conclusion can be no less than the probability of the conjunction of the premises. Thus, deductively valid arguments preserve not only *truth* from premises to conclusion, but also *degrees of evidential support*.³

Before you acquire the evidence of hypoxia, your premises about the distance of the journey and the amount of fuel remaining are highly probable given your evidence. Indeed, the probability of the conjunction of these premises is high enough that you can know the conclusion by deduction from these premises. Since your premises entail the conclusion, the probability of the conclusion is no less than the probability of the conjunction of the premises.

What changes when you acquire the new evidence that you are hypoxic? Now you have evidence that you are cognitively impaired in ways that dispose you to botch the reasoning from your premises to your conclusion. As we've seen, however, this doesn't undermine the entailment from premises to conclusion. Moreover, this doesn't affect the evidential probability of the premises themselves. The strength of your evidence for these premises remains unchanged. After all, hypoxia doesn't cause you to hallucinate the readings on your control panel or to invent new estimates of the distance to your destination. Rather, the point of the example is that it impairs your capacity to acquire knowledge by deduction.

Since the evidence of hypoxia doesn't change the evidential probability of the premises, it doesn't change the evidential probability of the conclusion either. Hence, acquiring this new evidence does nothing to reduce the probability of your conclusion. I conclude that the evidence of hypoxia is no evidential defeater at all, since it doesn't reduce the overall degree of evidential support for your conclusion. All evidential defeaters, whether or not they are evidence against a conclusion, reduce the degree to which your total evidence supports a conclusion. We can articulate this as a probabilistic constraint on evidential defeat:

A Probabilistic Constraint on Evidential Defeat: If d defeats the evidential support that e provides for h , then the probability that h given e and d is less than the probability that h given e alone.⁴

The higher-order evidence that you are hypoxic doesn't satisfy this probabilistic constraint, since the evidential probability that you have enough fuel remains unchanged when you receive this new evidence.

Some may be tempted to block this objection by rejecting the probabilistic conception of evidential support on which it depends. In response, however, this

³ See Skyrms (1966: Ch. 2) for a compelling articulation of this point. In the text, I assume *regularity* – that is, impossible scenarios have zero probability.

⁴ Compare Kotzen's thesis that defeaters are *credence lowering*: " D is a defeater for the evidence that E provides for H just in case $p(H | E \wedge D) < p(H | E)$ " (2019: 15).

comes at too high a cost. As we've seen, the probabilistic framework is designed to capture the important insight that arguments come in varying degrees of strength. Deductively valid arguments provide the strongest possible degree of support for a conclusion, since the premises entail the conclusion, rather than merely raising the probability of the conclusion. My objection exploits this intuitive point about the strength of deductive arguments, which is captured most naturally in a probabilistic framework.

To be clear, I am not questioning the intuitive datum that higher-order evidence of hypoxia destroys your knowledge that you have enough fuel. Moreover, I have no complaint about the practice of articulating this datum using the language of 'defeat' so long as we are careful to recognize that this is not an explanation of the datum but merely a restatement of the datum to be explained. The challenge that remains is to explain *why* this higher-order evidence destroys your knowledge. In this section, I've argued that the evidence of hypoxia doesn't undermine your knowledge by reducing your overall degree of evidential support for the conclusion that you have enough fuel. Some other explanation is needed.

2. Epistemic Conflicts

David Christensen (2007, 2010a) argues that higher-order evidence destroys knowledge by creating *epistemic conflicts*. Misleading higher-order evidence is "rationally toxic" in the sense that it forces you to violate one of the following epistemic ideals:

- (1) Respecting your first-order evidence.
- (2) Respecting your higher-order evidence.
- (3) Meta-coherence, i.e. coherently integrating your first-order beliefs with your higher-order beliefs.

In the hypoxia case, for example, you cannot respect all your evidence while coherently integrating your first-order and higher-order beliefs. Your first-order evidence supports the belief that you have enough fuel, while your higher-order evidence supports the higher-order belief that this first-order belief is probably based on bad reasoning. And yet this combination of beliefs seems dubiously coherent. As Christensen writes, "the rationality of first-order beliefs cannot in general be divorced from the rationality of certain second-order beliefs that bear on the epistemic status of those first-order beliefs" (2007: 18).

Christensen argues that the rationally optimal way of resolving this conflict is to violate the epistemic ideal of respecting your evidence. Rather than believing what your evidence supports – namely, that you have enough fuel – you should instead remain agnostic. Hence, epistemic rationality requires that you "bracket" your first-order evidence in the sense that you refrain from believing what it supports. On this view, your knowledge that you have enough fuel is destroyed when you acquire the evidence that you are hypoxic because it is no longer epistemically rational to believe what your evidence supports.

There is something puzzling about this proposal. How can epistemic rationality require you to refrain from believing what your evidence supports? According to *evidentialism* in epistemology, epistemic rationality is simply a matter of proportioning your beliefs to your evidence.⁵ On this view, epistemic rationality never requires or even permits you to “bracket” any of your evidence. On the contrary, it imposes a requirement of *total evidence*, which says you should always take all of your evidence into account in deciding what to believe.

Christensen’s proposal can be understood as a form of *bifurcationism* about epistemic rationality.⁶ On this view, the structural requirements of coherence, including meta-coherence, are distinct from and irreducible to the structural requirement of respecting your evidence. Moreover, these requirements can come into conflict when you have misleading higher-order evidence about your response to your first-order evidence, since your total evidence supports meta-incoherent beliefs. Christensen’s proposal is that the rationally optimal way of resolving these conflicting requirements is to maintain meta-coherence in response to your higher-order evidence by disrespecting your first-order evidence.

I argue elsewhere that we should prefer a *unified* conception of epistemic rationality, according to which the structural requirements of coherence are built into the structure of the evidential support relation (Smithies, forthcoming). For example, we can build in requirements of logical or probabilistic coherence by endorsing a probabilistic conception of the evidential support relation, according to which degrees of evidential support are evidential probabilities. Similarly, we can build in a meta-coherence requirement by endorsing structural constraints on higher-order probabilities, such as the following:

Probabilistic Accessibilism: Necessarily, if the evidential probability that p is n , then it is evidentially certain that the evidential probability that p is n .⁷

On this unified conception of epistemic rationality, there is no distinction to be drawn between the substantive requirement to respect your evidence and the structural requirement to be coherent. There is just one evidentialist requirement that incorporates both substantive and structural dimensions – that is, to

⁵ See Feldman and Conee (1985) for a classic defense of evidentialism.

⁶ Compare Worsnip (2018) for a similar view. See Smithies (forthcoming) for a detailed critical discussion of several different forms of bifurcationism, including those defended by Christensen and Worsnip.

⁷ Christensen (2010b) and Elga (2013) reject probabilistic accessibilism in favor of rational reflection principles, but these higher-order constraints are too weak to prohibit instances of epistemic akrasia in which you are certain that your credence is irrational, although you are agnostic about whether it is too high or too low. In Smithies (2019: Ch. 10), I defend probabilistic accessibilism against the objection that it is “immodest” by appealing to the distinction between ideal and non-ideal standards of epistemic rationality drawn in §5 below.

proportion your beliefs to your evidence in the sense that your beliefs cohere with substantive facts about your evidence in accordance with structural facts about the evidential support relation. These structural constraints on the evidential support relation guarantee that your evidence never supports an incoherent set of beliefs.

Occam's razor prohibits multiplying requirements of epistemic rationality beyond necessity. Why then might someone endorse bifurcationism? The usual answer is that bifurcationism is supported by reflection on examples. In the hypoxia case, for example, your evidence seems to support the following beliefs:

- (1) I have enough fuel.
- (2) But my belief that I have enough fuel is probably based on bad reasoning, since I am cognitively impaired as a result of hypoxia.
- (3) Therefore, my belief that I have enough fuel is probably not supported by good evidence.

And yet this combination of beliefs violates the meta-coherence requirement, since it is always irrational to hold a belief while also believing that it is neither based on nor supported by good evidence. Hence, the substantive requirement to respect your evidence seems to conflict in this case with the structural requirement of meta-coherence.

As we've seen, however, we can build the requirement of meta-coherence into the structure of the evidential support relation. Given probabilistic accessibilism, for example, you cannot have misleading higher-order evidence about what your evidence supports, since the facts about what your evidence supports are always made certain by your evidence. On this view, your evidence never supports a meta-incoherent combination of beliefs of the form, ' p and my evidence probably doesn't support p '. Hence, respecting your evidence guarantees that you are also meta-coherent.

Now, of course, you can have misleading higher-order evidence about your *response* to your evidence, although you can never have misleading higher-order evidence about what your evidence *supports*. This is because facts about your response to your evidence, unlike facts about what your evidence supports, are not made certain by your evidence. As a general rule, the argument from (2) to (3) is inductively strong, since beliefs based on bad reasoning are not usually supported by good evidence. In your own case, however, the inference from (2) to (3) is blocked when your higher-order evidence is misleading. In such cases, your evidence supports believing (1) and (2) but not (3).

Moreover, believing (1) and (2) is not incoherent in the same way as believing (1) and (3). It's always irrational to believe that p while believing that your evidence doesn't support p . This combination of beliefs is self-defeating: by your own lights, you should abandon your belief that p , since you think it is unsupported by your evidence. In contrast, it's not always irrational to believe that p while believing that your belief is not properly based on supporting evidence. Indeed, this can be a perfectly rational response to misleading evidence that you

believe the right thing for the wrong reasons. It is not self-defeating in such cases to conclude that your belief is supported by good evidence, although it is not properly based on good evidence.

To illustrate the point, consider a practical analogy. Suppose a wealthy philanthropist receives public acclaim for donating large sums of money to charity. He knows he is doing the right thing, but he suspects he is doing it for the wrong reasons because he has misleading evidence that he is selfishly motivated. So long as this evidence is misleading, his action may be rationally responsive to good moral reasons. The same applies to someone who knows that he believes what his evidence supports, although he suspects his belief is held for the wrong reasons because he has misleading evidence that he is motivated by wishful thinking. Once again, his belief may be rationally responsive to good evidence so long as this higher-order evidence is misleading. In such cases, there is nothing self-defeating about doubting the reasoning on which your beliefs and actions are based.

Admittedly, there is something unusual about an evidential situation that supports the following line of argument:

- (4) It's certain that my evidence supports the conclusion that I have enough fuel.
- (5) But my belief that I have enough fuel is probably based on bad reasoning.
- (6) So I probably got lucky, since bad reasoning led me to form a belief that is supported by the evidence.

In the absence of strong evidence, it's irrational to believe in lucky coincidences. After all, the prior probability of such a coincidence is very low. In the presence of strong evidence, however, the posterior probability of a coincidence may be high. There is nothing in principle to rule out the possibility that your evidence supports the conclusion that a lucky coincidence has occurred. Otherwise, no one could find out when they win the lottery. The hypoxia example is another case in point, since your evidence supports the conclusion that you got lucky.

Does this license dogmatism in response to higher-order evidence? Consider our steadfast pilot who maintains his belief that he has enough fuel in the face of the higher-order evidence that he is hypoxic. Suppose that instead of simply ignoring this higher-order evidence, he concludes that he must have got lucky, since his belief is supported by evidence although it is probably based on bad reasoning. Intuitively, this compounds his irrationality. It was already irrational for him to retain the first-order belief that he has enough fuel and the problem is exacerbated when he doubles down by retaining the higher-order belief that his evidence supports this conclusion. Our steadfast pilot maintains meta-coherence at the cost of both first-order and higher-order dogmatism.

I agree the steadfast pilot is irrational, but this just shows that our original problem arises at multiple levels. At level one, the problem is to explain why the pilot cannot rationally believe that he has enough fuel when his evidence supports this conclusion. At the level two, the problem is to explain why the pilot cannot

rationally believe that his evidence supports this conclusion when his evidence also supports this higher-order conclusion. And so on as we ascend the hierarchy. I'll propose my own solution to this problem in due course. My goal in this section is merely to argue that we cannot solve it by arguing that the pilot's evidence supports incoherent beliefs. There are general theoretical reasons to doubt that your evidence can ever support incoherent beliefs and there are more specific reasons to doubt that the pilot's evidence supports incoherent beliefs. We need to look elsewhere to explain why it is irrational for the steadfast pilot to retain beliefs that are nevertheless supported by his evidence.

3. Unreasonable Knowledge

Maria Lasonen-Aarnio (2010, 2020) argues that the steadfast pilot has *unreasonable knowledge*.⁸ On this view, the pilot can retain his knowledge that he has enough fuel even in the face of misleading higher-order evidence that he is hypoxic. The problem is that his belief is unreasonable because it manifests a more general disposition to be unresponsive to evidence in other cases. And yet this needn't undermine his knowledge so long as he responds appropriately to all the evidence that he actually has. Thus, reasonable belief is not necessary for knowledge.

What is it for a belief to be reasonable? Lasonen-Aarnio writes:

Reasonableness is at least largely a matter of managing one's beliefs through the adoption of policies that are generally knowledge conducive, thereby manifesting dispositions to know and avoid false belief across a wide range of normal cases. (2010: 2)

The steadfast pilot is unreasonable because he fails to manifest dispositions that are generally conducive to knowledge – that is, dispositions to know and avoid false belief across a wide range of normal cases. In particular, he is disposed to ignore higher-order evidence that his beliefs are based on bad reasoning not only in the “good case” in which his higher-order evidence is misleading but also in the “bad case” in which it is accurate. Since the actual case is a good case, his belief is not only true, but also supported by his evidence. Nevertheless, it is held dogmatically in a way that disposes him to retain beliefs in bad cases that are both false and unsupported by evidence.

In short, the steadfast pilot is unreasonable because he manifests a general disposition that leads him astray in other cases. And yet this leaves his knowledge intact so long as the disposition doesn't lead him astray in this case. He retains his knowledge because his true belief remains as safe from error, and as responsive to his first-order evidence, as it was before he acquired the higher-order evidence of hypoxia. Since he had knowledge beforehand, he retains his knowledge in the face of his new higher-order evidence, despite the fact that his belief is unreasonable.

⁸ Lasonen-Aarnio (2010) focuses on standard examples of undercutting defeat, but she extends her proposal to higher-order evidence in Lasonen-Aarnio (2020).

It's worth noting that the same proposal applies at multiple levels. Presumably, the steadfast pilot can retain not only his first-order knowledge that he has enough fuel, but also his higher-order knowledge that his evidence supports this conclusion. Moreover, given the misleading evidence that he is hypoxic, he can rationally conclude from these known premises that he probably got lucky, since his cognitive impairment didn't lead him astray on this occasion. The problem with his first-order and higher-order beliefs is that they are unreasonable because they manifest more general dispositions to go awry in bad cases in which his higher-order evidence is accurate. But this has no tendency to undermine either the pilot's first-order knowledge or his higher-order knowledge.

I'll now raise three objections to this proposal. First, it fails to vindicate all our intuitions about the hypoxia case. It vindicates the intuition that the steadfast pilot is unreasonable, but not the intuition that he loses knowledge. Is it plausible that the steadfast pilot retains his first-order knowledge that he has enough fuel as well as his higher-order knowledge that his evidence supports this conclusion? Is it plausible that he can use this knowledge, together with his misleading evidence that he is hypoxic, to infer that he probably got lucky on this occasion? To many, myself included, these verdicts are strongly counterintuitive.

Lasonen-Aarnio proposes to explain away conflicting intuitions by appeal to the error theory that we tend to conflate reasonableness and knowledge. I agree that our intuitions about this case, and many others, are guided by the implicit assumption that only reasonable beliefs can be knowledge. This explains why we find it so natural to make the inference from the premise that someone's belief is unreasonable to the conclusion that they lack knowledge. Given the intuitive plausibility of this assumption, however, why suppose it is mistaken? To my mind, the error theory should be regarded as a last resort: all else being equal, we should prefer an epistemological theory that vindicates our intuitive reactions.

Second, there are general theoretical grounds for doubting that you can acquire knowledge by manifesting unreasonable dispositions. To know that p , it's not enough that you truly believe that p ; your belief must be *reliable* in the sense that it manifests a more general disposition to have true beliefs. Similarly, to rationally or justifiably believe that p , it's not enough that your belief is supported by good evidence; it must be *properly based* on good evidence in the sense that it manifests a more general disposition to have beliefs that are supported by good evidence. Hence, knowledge requires manifesting good dispositions that are reliably responsive to evidence and truth. When your beliefs manifest unreasonable dispositions, however, they are not reliable enough to constitute knowledge. Presumably, this is why it remains so intuitively plausible that only reasonable beliefs can be knowledge. This is not just a brute intuition with no theoretical support. It is supported by the theoretical consideration that knowledge must be reliably responsive to evidence and truth.

Lasonen-Aarnio sympathetically discusses the simple externalist view that knowledge is true belief that is *safe from error* in the sense that it couldn't easily have been false. But the case of mathematical knowledge suggests that this view is

too simple: I cannot know that Fermat's Last Theorem is true by means of wishful thinking, although there is no danger that my belief is false, since its content is necessarily true. Arguably, this is because my safe belief is based on manifesting an *unsafe disposition*: one that could easily yield false beliefs in other cases. This supports the claim that a belief is knowledge only if it is reasonable in the sense that it manifests a more generally reliable disposition.⁹

Third, there are theoretical costs involved in rejecting the principle that only reasonable beliefs can be knowledge. Knowledge is valuable. If we allow that knowledge can be unreasonable, however, then we risk devaluing knowledge. So, for example, knowledge is often thought to set a normative standard of correctness for belief and action: if you know that p , then you are right to believe that p , and to act on the premise that p , and so you cannot legitimately be blamed for doing so. And yet the steadfast pilot is blameworthy for acting on the premise that he has enough fuel: his mother can be justly angry that he acted so recklessly. As Lasonen-Aarnio acknowledges, "Subjects who retain knowledge in defeat cases are genuinely criticisable" (2010: 15). The problem is that we cannot maintain that the steadfast pilot retains his knowledge, while also acknowledging his culpability, unless we abandon these plausible connections between knowledge, permission, and blameworthiness. But this devalues knowledge in ways that should be unattractive to anyone, especially to proponents of knowledge-first epistemology.

In conclusion, we need to explain not only how evidence of hypoxia can make your beliefs unreasonable, but also how it can destroy your knowledge. Although Lasonen-Aarnio's proposal fails to explain this datum, it contains an important grain of truth, which will figure prominently in the discussion to follow. To explain how higher-order evidence destroys your knowledge, we need to invoke facts about the reliability of your doxastic dispositions.

4. Improper Basing

My own explanation of how higher-order evidence destroys first-order knowledge appeals to *improper basing*.¹⁰ On this view, the higher-order evidence that you are hypoxic doesn't defeat your first-order evidence that you have enough fuel, but merely prevents you from properly basing your beliefs on this evidence. This explains why you lose your knowledge that you have enough fuel when you acquire the higher-order evidence that you are hypoxic. Only justified beliefs can be knowledge and a belief is justified only if it is properly based on good evidence.

⁹ See Sosa (2003) for a defense of this claim. In Smithies (2019: Ch. 11), I argue for similar conclusions about justification based on the problem of the speckled hen: justified belief requires exercising a disposition that is reliably responsive to what your evidence supports.

¹⁰ Compare Smithies (2015; 2019: Ch. 10). Van Wietmarschen (2013) adopts a similar line on peer disagreement, although his explanation of how higher-order evidence undermines proper basing diverges from mine.

We can articulate the point in terms of the familiar distinction between *propositional* and *doxastic* senses of epistemic justification, rationality, or warrant.¹¹ A belief is propositionally justified when its propositional content is supported by good evidence, whereas a belief is doxastically justified when the belief is held in a way that is properly based on good evidence. Thus, *proper basing* is the relation between a belief and a body of supporting evidence that converts propositional justification into doxastic justification. The higher-order evidence that you are hypoxic doesn't undermine your propositional justification to believe that you have enough fuel by defeating your evidence for this conclusion. Rather, it undermines your doxastic justification by preventing you from properly basing your belief on this evidence. In a slogan, higher-order evidence is a *doxastic defeater*, rather than a *propositional defeater*.

What does it take for a belief to be properly based on good evidence? It's not enough that my belief is based on good evidence that happens to support my beliefs. My belief must also manifest a more general disposition that is reliably sensitive to differences in what my evidence supports. If I'm disposed to retain my belief even if my evidence changes in ways that no longer support the belief, then it is not properly based on the evidence. My belief is properly based on the evidence only if it manifests a more general disposition to believe what the evidence supports.¹²

Dogmatic beliefs violate this condition. Consider my belief that drinking red wine is good for my health. Let's assume that while the evidence for this claim is somewhat mixed, the evidence in its favor outweighs the evidence against, and that the supporting evidence is strong enough to justify belief. Although I'm aware of all this evidence, the problem is that I hold my belief dogmatically in a way that makes me relatively insensitive to changes in what my evidence supports. I am not disposed to respond to changes in what my evidence supports with corresponding changes in what I believe. I will remain unmoved, for example, if I learn about new studies casting doubt on the health benefits of drinking red wine. Intuitively, I don't know that drinking wine is good for me, even if my belief is true and based on good evidence. The problem is that my belief is not *properly based* on my evidence, since it manifests a disposition that is insufficiently sensitive to changes in what my evidence supports.

I suggest that the same is true of the steadfast pilot. He doesn't know that he has enough fuel, despite the fact that his belief is true and based on good

¹¹ This distinction is usually traced back to Firth (1978), although it is now ubiquitous in the literature on epistemic justification.

¹² In Smithies (2019: Ch. 11), I use the problem of the speckled hen to motivate this constraint on proper basing. More specifically, I argue that a belief is properly based on supporting evidence only if it manifests a more general disposition to form beliefs that are safe from the absence of evidential support. On this view, doxastic justification requires safety from the absence of propositional justification just as knowledge requires safety from error. I don't rely on this proposal here.

evidence, because it is not *properly based* on his evidence. The intuitive problem is that his belief is held dogmatically, which means it is not sufficiently sensitive to changes in what his evidence supports. We can see this by comparing the “good case” in which his higher-order evidence is misleading with the “bad case” in which his higher-order evidence is accurate.

Let’s set aside modally remote cases in which the steadfast pilot is actually suffering from hypoxia, since this dramatically alters his capacity for reasoning. Instead, let’s consider more quotidian cases in which his reasoning capacities are held constant. Since his capacities are fallible, there are cases much closer to home in which he makes routine errors in calculation through the defective exercise of those very capacities. Let’s consider a bad case in which he makes a routine error and his co-pilot points out the mistake. If he is disposed to remain steadfast in the face of evidence that he is hypoxic, then he will be equally disposed to ignore his co-pilot and stick to his guns. Moreover, both cases manifest the same disposition to dogmatically retain beliefs in the face of higher-order evidence that those beliefs are based on bad reasoning.

As we saw in §3, Lasonen-Aarnio makes similar points in arguing that the steadfast pilot is unreasonable in retaining his belief that he has enough fuel. In contrast, my goal here is to argue that his belief is doxastically unjustified, and hence not a case of knowledge, since it is not properly based on supporting evidence. In the good case, his belief is true and supported by evidence; but in the bad case, it is false and unsupported by evidence. So, even in the good case, the steadfast pilot doesn’t have knowledge, since his belief manifests a disposition that is not sufficiently sensitive to changes in what his evidence supports. After all, there are close cases in which exercising the same disposition leads him to hold false beliefs in the absence of evidential support.

We can apply the same reasoning one level up to explain why the steadfast pilot cannot have *higher-order* knowledge or justified belief about his own epistemic situation. Suppose he believes not only that he has enough fuel but also that his evidence supports this conclusion. Although this higher-order belief is true and supported by conclusive evidence, it is not properly based on this evidence. After all, the steadfast pilot is disposed to believe exactly the same thing in the bad case in which his higher-order belief is false and unsupported by his evidence. So, even in the good case, he doesn’t know that his evidence supports his conclusion, since his belief manifests a disposition that is not sufficiently sensitive to changes in what his evidence supports.

I am not claiming that that the steadfast pilot *loses* knowledge when he acquires the higher-order evidence that he is hypoxic. Indeed, there is something deeply puzzling about this suggestion.¹³ What changes when he acquires the evidence that he is hypoxic? There is no relevant change in what his evidence supports, since it continues to support the conclusion that he has enough fuel. Moreover, there is no relevant change in his responsiveness to his evidence, since

¹³ Lasonen-Aarnio (2010: 3-8) raises a version of this puzzle.

acquiring the higher-order evidence makes no difference to the basis on which his belief is held. But then how can it be that acquiring this higher-order evidence destroys his knowledge?

My answer is that the steadfast pilot doesn't *lose* knowledge, since he never had knowledge to begin with. You cannot acquire knowledge by manifesting dogmatic dispositions that are insensitive to changes in what your evidence supports. In contrast, the *conciliatory pilot* loses knowledge when he acquires the misleading higher-order evidence because he responds by abandoning his belief. There can be no knowledge without belief. More importantly, however, knowledge sometimes requires being disposed to abandon belief in response to new evidence. The conciliatory pilot has knowledge before acquiring the evidence that he is hypoxic only because he is disposed to respond to this new evidence by abandoning his belief. The conciliatory pilot is more reliably sensitive to evidence than the steadfast pilot because he is not disposed in bad cases to ignore accurate higher-order evidence that he has made a mistake.

I don't claim that it's impossible in principle to retain knowledge in the face of misleading higher-order evidence, but merely that it's impossible in practice for creatures like us. I see no reason to rule out the metaphysical possibility of an ideally rational agent who is perfectly sensitive to what her evidence supports. Because she is perfectly sensitive to her evidence, she can remain steadfast in good cases without running the risk of remaining steadfast in bad cases in which her reasoning dispositions are held constant.¹⁴ And yet this is beyond the capacity of any normal human agent, since we are only imperfectly sensitive to our evidence. Any human who remains steadfast in good cases thereby manifests some disposition to remain steadfast in bad cases too.

There is a more general moral to be drawn here. There are cases in which your evidence supports a conclusion, and gives you propositional justification to believe it, although you are psychologically incapable of forming a doxastically justified belief that is properly based on your evidence. We should therefore reject the doxastic constraint on propositional justification stated below:

The Doxastic Constraint: Necessarily, you have propositional justification to believe that p only if you have the psychological capacity to believe that p in a way that is doxastically justified.¹⁵

There are independent reasons to reject the doxastic constraint. Suppose you're given a reason-distorting drug that renders you temporarily incapable of responding rationally to your evidence. Whatever doxastic attitudes you adopt, they are guaranteed to be doxastically unjustified. It seems absurd to suppose that

¹⁴ I don't claim that an ideally rational agent must be immune from hypoxia. But we can ignore bad cases in which she suffers from hypoxia, since this changes her reasoning dispositions in ways that make them less than ideally rational.

¹⁵ Proponents of the Doxastic Constraint include Goldman (1979) and Turri (2010).

merely ingesting this drug undermines your propositional justification to adopt any doxastic attitude at all. There is always some doxastic attitude that you have propositional justification to hold towards any given proposition. If your evidence is not strong enough to justify either belief or disbelief, then you have justification to remain agnostic by default. Even so, there is no guarantee that you are always psychologically capable of forming a doxastic attitude that is properly based on what your evidence supports. After all, you might just have ingested a reason-distorting drug.

The natural thing to say about the reason-distorting drug is that it impairs your epistemic rationality. It prevents you from properly basing your beliefs on your evidence and thereby converting your propositional justification into doxastic justification. If there are doxastic constraints on propositional justification, however, then we cannot say this. We must say instead that ingesting the drug somehow changes what you have propositional justification to believe. But this is hard to reconcile with the evidentialist thesis that you have propositional justification to believe whatever is sufficiently supported by your evidence, since ingesting the drug doesn't change what your evidence is or what it supports. Moreover, it loses sight of the intuitive idea that ingesting the drug compromises your epistemic rationality by preventing you from responding appropriately to your evidence.

On a plausible version of evidentialism, there are no doxastic constraints on propositional justification. What your evidence supports is one thing, but it's another issue entirely whether or not you're capable of responding appropriately by forming beliefs that are not only supported by your evidence, but also properly based on your evidence. One theoretical cost of conflating these questions is to obscure the epistemic function of higher-order evidence.

5. Non-Ideal Rationality

Any plausible account of the epistemic function of higher-order evidence must explain two intuitive data points about the hypoxia example:

- (1) *The Negative Datum*: You cannot know or rationally believe that you have enough fuel given higher-order evidence that you're hypoxic
- (2) *The Positive Datum*: You are rationally required to withhold belief that you have enough fuel given higher-order evidence that you're hypoxic.

I explained the negative datum in §4, but this is not yet to explain the positive datum. After all, a reason-distorting drug might prevent you from rationally believing what your evidence supports without thereby imposing any rational requirement to refrain from believing what your evidence supports. Indeed, there is something puzzling about the very idea of such a requirement. According to evidentialism, epistemic rationality is simply a matter of proportioning your beliefs to your evidence. So how can epistemic rationality ever require you to refrain from believing what your evidence supports?

To answer this question, we need a distinction between *ideal* and *non-ideal* standards of epistemic rationality.¹⁶ By ideal standards, epistemic rationality always requires respecting your evidence. This is an epistemic ideal that sometimes falls beyond our limited human capacities: we are not always capable of rationally following our evidence where it leads. Moreover, we know this – at any rate, we should know this – since we all have compelling evidence of our own cognitive limitations. By non-ideal standards, in contrast, epistemic rationality sometimes requires responding to such evidence by adopting policies that diverge from the epistemic ideal. In particular, we are sometimes required by non-ideal standards to “bracket” our first-order evidence when we have higher-order evidence that we cannot reliably follow our first-order evidence where it leads.

This is what happens in the hypoxia example. When you acquire the higher-order evidence that you cannot respond rationally to your first-order evidence, you are required to “bracket” this evidence and refrain from believing what it supports – namely, that you have enough fuel. Given evidence of cognitive impairment, the sensible strategy is to become agnostic, rather than trying to believe what your evidence supports. On this view, the epistemic function of higher-order evidence is not to defeat your first-order evidence and thereby to affect which conclusions are supported by your total evidence. Rather, it determines which response to your evidence is required by non-ideal standards of epistemic rationality.

This is an instance of a much more general point. It doesn’t always make sense to try to do what you know would be best if only you succeed in doing it. After all, you sometimes know – or have good evidence – that your attempt to do the best thing may not succeed. In such cases, trying to do the best thing often risks a worse outcome than would be achieved by settling for second best. Hence, what is best by ideal standards is not always what is best by non-ideal standards that take into account your evidence about your own limitations.

We can articulate this general point more precisely within the framework of *rule consequentialism*, which evaluates rules by their expected consequences. We can evaluate rules in a way that is sensitive to the distinction between *following* a rule and merely *trying* to follow a rule. Following a rule is a kind of achievement: merely trying to follow the rule doesn’t guarantee that you will succeed. When you have evidence that you might fail, the expected consequences of trying to follow a rule can diverge from those of successfully following the rule. In such cases, the best rule to follow is not always the best rule to try to follow.¹⁷

Consider Professor Procrastinate who is invited to review a book.¹⁸ He has three options available, which are listed below in rank order from best to worst:

- (1) Accept the invitation and complete the review on time.
- (2) Decline the invitation.

¹⁶ Compare Smithies (2015; 2019: Ch. 10; forthcoming) for this distinction.

¹⁷ Compare Lasonen-Aarnio (2010: 14-15) and Schoenfield (2015: 650-3).

¹⁸ The example is from Jackson and Pargetter (1986).

- (3) Accept the invitation and fail to complete the review on time.

Which option should he take? There is no single answer to this question, since deontic modals in ordinary language are highly context-sensitive. There is a sense in which he should accept the invitation and complete the review on time, since he knows this is the best possible outcome. But he also knows (or has good evidence) that he won't complete the review on time: if he tries to achieve the best outcome, then he is more likely to bring about the worst outcome.¹⁹ Hence, there is a sense in which he should decline the invitation. Rather than trying to do what is best, it makes more sense to settle for second best. By ideal standards, he should accept the invitation and complete the review on time, whereas by non-ideal standards, he should decline.

We can capture this distinction in the framework of rule-consequentialism. Rule (1) is the best rule to follow, but it is not the best rule to try to follow. After all, the expected consequence of trying to follow Rule (1) is that he will end up following Rule (3) instead. The expected consequence of trying to follow Rule (2), however, is that he will succeed. And there is greater expected value in following Rule (2) than Rule (3). Hence, there is greater expected value in trying to follow Rule (2) than Rule (1), despite the fact that there is greater expected value in successfully following Rule (1) than Rule (2). This is why ideal standards require following Rule (1), although non-ideal standards require following Rule (2).

When we're evaluating epistemic rules, we're concerned solely with their expected epistemic consequences. And when we're evaluating them for epistemic rationality, we're specifically concerned with their expected consequences for how well you succeed in proportioning your beliefs to your evidence.²⁰ From an evidentialist perspective, the best rule to follow is *the evidentialist rule*, "Always proportion your beliefs to your evidence!" By ideal standards, epistemic rationality always requires following the evidentialist rule.

Since we are not perfectly rational agents, however, we are not always capable of following the evidentialist rule. Moreover, this is not always the best rule to try to follow when you have evidence that you may fail. There may be greater expected epistemic value in trying to follow some alternative rule, since you are likely to manifest more evidence-sensitive dispositions by adopting this alternative strategy. In such cases, trying to follow the evidentialist rule is a kind of self-sabotage: it's a counterproductive strategy for maximizing your responsiveness

¹⁹ As far as I can discern, it makes no difference whether his evidence is misleading. Perhaps there are contexts in which our intuitions about what people ought to do track their actual tendencies, rather than their evidence about them, but there are also contexts in which it is their evidence that matters.

²⁰ In contrast, Lasonen-Aarnio (2010) is concerned with expected epistemic consequences for knowledge, while Schoenfield (2015) is concerned with expected accuracy, whereas I am concerned with expected epistemic rationality construed as responsiveness to evidence.

to evidence. This is why you are sometimes required by non-ideal standards of epistemic rationality to refrain from believing what your evidence supports.

Now let's apply this distinction to the hypoxia example. When you receive the higher-order evidence that you are hypoxic, you have three options:

- (1) *Steadfastness*: Maintain your first-order belief that you have enough fuel and also your higher-order belief that your evidence supports this conclusion.
- (2) *Level Splitting*: Maintain your first-order belief that you have enough fuel, but abandon your higher-order belief that your evidence supports this conclusion.
- (3) *Conciliation*: Abandon your first-order belief that you have enough fuel and also your higher-order belief that your evidence supports this conclusion.

Which option should you take? Again, there is no single answer this question. By ideal standards, you should follow *Steadfastness*, since this is the best rule to follow: the expected consequence of successfully following this rule is that your beliefs are proportioned to your evidence. By non-ideal standards, however, you should follow *Conciliation*, since this is the best rule to try to follow. As I'll explain, there is greater expected epistemic value in trying to follow *Conciliation* than either *Steadfastness* or *Level Splitting*.²¹

The expected consequence of trying to follow *Steadfastness* is that you will be like the *steadfast pilot*.²² As we saw in §4, the steadfast pilot cannot know or rationally believe that he has enough fuel whether or not he has higher-order evidence that he is hypoxic. This is because his belief is held dogmatically in a way that is insufficiently sensitive to changes in what his evidence supports. He is disposed to maintain his belief when he acquires the higher-order evidence that it is based on bad reasoning. And he is disposed to maintain his belief not only in the good case in which his higher-order evidence is misleading, but also in the bad case in which his higher-order evidence is accurate. Dogmatic beliefs of this kind are not rational enough to constitute knowledge.

In contrast, the expected consequence of trying to follow *Conciliation* is that you will be like the *conciliatory pilot*. The conciliatory pilot rationally believes and knows that he has enough fuel before acquiring the higher-order evidence that he is hypoxic, so long as he is disposed to become agnostic when he acquires this higher-order evidence. Hence, there is greater expected epistemic value in trying to follow *Conciliation* than *Steadfastness*: this strategy increases your expected degree of epistemic rationality by making your beliefs more sensitive to changes in

²¹ Schoenfield (2015: 652) gives a different argument for the same conclusion.

²² Similarly, the expected consequence of trying to follow *Level Splitting* is that you will be like the steadfast pilot at level one and the conciliatory pilot at level two.

what your evidence supports. This is why you are required by non-ideal standards of epistemic rationality to become agnostic when you acquire the higher-order evidence that you are hypoxic. In this way, we can explain the positive datum as well as the negative datum.

In summary, the epistemic function of higher-order evidence is not only to prevent you from responding properly to your evidence in the way that is required by ideal standards of epistemic rationality. It also affects which response to your evidence is required by non-ideal standards of epistemic rationality that are sensitive to your evidence about your cognitive limitations.

6. Conclusions

Let me conclude by summarizing the reasons for preferring my account of the epistemic function of higher-order evidence to the alternatives discussed earlier in this chapter.

First, my view explains the intuitive datum that you lose knowledge when you acquire the higher-order evidence that you are hypoxic. In this respect, it is preferable to the *unreasonable knowledge* view. What this view gets right is that the steadfast pilot manifests bad dispositions when he retains his belief in the face of misleading higher-order evidence. What it gets wrong is that manifesting these bad dispositions is compatible with knowledge. The steadfast pilot doesn't have knowledge because his dispositions are not sufficiently sensitive to changes in what his evidence supports. It's possible in principle for ideally rational agents to retain knowledge in the face of misleading higher-order evidence, but this is impossible in practice for non-ideal agents like us.

Second, my view explains this intuitive datum without distorting the structure of the evidential support relation in ways that compromise the objective constraints imposed by logic and probability theory. In this respect, it is preferable to the *evidential defeat* view. What this view gets right is that your knowledge is destroyed when you acquire the higher-order evidence that you are hypoxic. What it gets wrong is that this is to be explained by a change in what your evidence supports. My view explains how you can lose your knowledge of a conclusion without losing your evidential support for that conclusion. Misleading higher-order evidence can prevent you from properly basing your beliefs on what your total evidence supports.

Third, my view explains this intuitive datum without any bifurcation between substantive and structural requirements of epistemic rationality. In this respect, it is preferable to the *epistemic conflict* view. What this view gets right is that epistemic rationality somehow requires "bracketing" your first-order evidence when you have misleading higher-order evidence. What it gets wrong is that this is because your total evidence supports incoherent beliefs. On the contrary, it is because misleading higher-order evidence prevents non-ideal agents from properly basing their beliefs on their evidence. As a result, they are required by non-ideal standards of epistemic rationality to manage their cognitive limitations by adopting policies that deviate from the epistemic ideal of respecting the evidence.

The “rationally toxic” nature of higher-order evidence is best explained in terms of a distinction between ideal and non-ideal requirements of epistemic rationality, rather than a distinction between substantive and structural requirements of epistemic rationality.

In conclusion, my proposal accommodates the intuitive data about the epistemic function of higher-order evidence with minimal theoretical mutilation. We can explain why you should conciliate in response to misleading higher-order evidence without abandoning evidentialism or compromising the objective logical and probabilistic constraints on the evidential support relation. The key point is that we are not always capable of rationally believing what our evidence supports. Moreover, we sometimes know or have misleading evidence that we are in this unfortunate predicament. In such cases, we should adopt strategies for managing our epistemic limitations that deviate from the epistemic ideal of believing what our evidence supports. This is one instance of the more general point that you shouldn’t always try to do what is best when your efforts are likely to backfire. Sometimes, you should settle for second-best.

References

- Christensen, David. 2007. “Does Murphy’s Law Apply in Epistemology? Self-Doubt and Rational Ideals.” *Oxford Studies in Epistemology* 2: 3-31.
- Christensen, David. 2010a. “Higher-Order Evidence.” *Philosophy and Phenomenological Research* 81.1: 185-215.
- Christensen, David. 2010b. “Rational Reflection.” *Philosophical Perspectives* 24: 121-140.
- Elga, Adam. 2008. “Lucky to Be Rational.” *Bellingham Summer Philosophy Conference*, unpublished manuscript.
- Elga, Adam. 2013. “The Puzzle of the Unmarked Clock and the New Rational Reflection Principle.” *Philosophical Studies* 164.1: 127-139.
- Feldman, Richard. 2005. “Respecting the Evidence.” *Philosophical Perspectives* 19: 95-119.
- Feldman, Richard and Earl Conee. 1985. “Evidentialism.” *Philosophical Studies* 48.1: 15-34.
- Firth, Roderick. 1978. “Are Epistemic Concepts Reducible to Ethical Concepts?” In *Values and Morals: Essays in Honor of William Frankena, Charles Stevenson, and Richard Brandt*, edited by Alvin Goldman and Jaegwon Kim, 215-229. Dordrecht: Kluwer.
- Goldman, Alvin. 1979. “What is Justified Belief?” In *Justification and Knowledge*, edited by George Pappas, 1-25. Dordrecht: Reidel.
- Hawthorne, John and Jason Stanley. 2008. “Knowledge and Action.” *Journal of Philosophy* 105.10: 571-590.
- Jackson, Frank and Robert Pargetter. 1986. “Oughts, Options, and Actualism.” *Philosophical Review* 95.2: 233-255.

- Kotzen, Matthew. 2019. "A Formal Account of Evidential Defeat." In *Themes from Klein: Knowledge, Justification, and Skepticism*, edited by Branden Fitelson, Rodrigo Borges, and Cherie Braden. Cham: Synthese Library.
- Lasonen-Aarnio, Maria. 2010. "Unreasonable Knowledge." *Philosophical Perspectives* 24: 1-21
- Lasonen-Aarnio, Maria. 2020. "Enkrasia or Evidentialism?" Learning to Love Mismatch. *Philosophical Studies* 177: 597-632.
- Pollock, John. 1986. *Contemporary Theories of Knowledge*. London: Hutchinson.
- Schoenfield, Miriam. 2015. "Bridging Rationality and Accuracy." *Journal of Philosophy* 112.12: 633-657.
- Skyrms, Brian. 1966. *Choice and Chance: An Introduction to Inductive Logic*. Belmont, CA: Dickenson.
- Smithies, Declan. 2015. "Ideal Rationality and Logical Omniscience." *Synthese* 192.9: 2769-93.
- Smithies, Declan. 2019. *The Epistemic Role of Consciousness*. Oxford University Press.
- Smithies, Declan. Forthcoming. "The Unity of Evidence and Coherence." In *Epistemic Dilemmas*, edited by Nick Hughes. Oxford University Press.
- Sosa, Ernest. 2003. "Beyond Internal Foundations to External Virtues." In *Epistemic Justification: Internalism vs. Externalism, Foundations vs. Virtues*, by Laurence Bonjour and Ernest Sosa. Oxford: Blackwell.
- Turri, John. 2010. "On the Relationship Between Propositional and Doxastic Justification." *Philosophy and Phenomenological Research* 80.2: 312-326.
- van Wietmarschen, Han. 2013. "Peer Disagreement, Evidence, and Well-Foundedness." *Philosophical Review* 122.3: 395-425.
- Worsnip, Alex. 2018. "The Conflict of Evidence and Coherence." *Philosophy and Phenomenological Research* 96.1: 3-44.