The Epistemic Function of Higher-Order Evidence

Declan Smithies
In *Propositional and Doxastic Justification: New Perspectives in Epistemology*,
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 travel 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 tend to remain undetectable to the victim. You know about the dangers of hypoxia: pilots have crashed and died because 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 the evidence for hypoxia is misleading. This evidence indicates that he is cognitively impaired in ways that dispose him to make basic errors in calculation. It seems reckless for him to ignore the possibility that he has miscalculated in deciding to stick with his original plan. Of course, we know there is no risk of hypoxia and so we can rest assured that the pilot will arrive safely. But the pilot himself doesn't know this. Anyone who routinely ignores evidence of hypoxia in this way 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: if it's rationally permissible to believe that *p*, then it's rationally permissible to use *p* as a premise in reasoning.[2] When you acquire the new evidence that you're hypoxic, it's rationally impermissible to reason from the premise that you have enough fuel because it's rationally impermissible to believe

---

[1] This example is adapted from Elga (2008; 2013) and Christensen (2010a).
[2] Some (e.g. Jackson 2019) reject this principle for high-stakes cases, but this feature of the example is incidental, since the same intuitive reactions are plausible in low-stakes cases too, including Christensen's (2010a, pp. 186-7) cases of peer disagreement and reason-distorting drugs.

you have enough fuel. Moreover, this means you cannot *know* you have enough fuel, since it's rationally permissible to believe that *p*, and hence to reason from the premise that *p*, whenever you know that *p*.[3] Previously, however, you did know 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 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 appropriately responsive to your evidence.[4] Suppose you know that *p* on the basis of evidence *e*, but you subsequently acquire misleading higher-order evidence *h* that your belief is unresponsive 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 evidence no longer supports the conclusion that you have enough fuel.

The challenge is to explain why the higher-order evidence that you are hypoxic defeats your first-order evidence that you have enough fuel. After all, your

---

[3] See Fantl and McGrath (2009, p. 66). Hawthorne and Stanley (2008, p.578) endorse the biconditional that it's rationally permissible to reason from the premise that *p* if and only if you know that *p*, but this is more controversial.

[4] I reject the usual definition of higher-order evidence as evidence about what your evidence supports, since I deny that you can have misleading evidence about what your evidence is or what it supports (see section 2).

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 the 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 the overall degree of evidential support for a conclusion without giving you evidence against the conclusion. As Pollock writes, "Such defeaters attack the connection between the reason and the conclusion rather than attacking the conclusion itself" (1986, p. 196).

Consider a textbook example: the fact that the wall looks red is evidence that it is red, although this evidential connection can be undercut by background evidence that the wall is bathed in red light. Of course, this is no evidence *against* the conclusion that the wall is red, since the lighting alone makes it no more or less probable that the wall is red. As I'll explain, however, the background evidence reduces the degree to which your sensory evidence supports the 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 makes it highly probable that the wall is red. Given the background evidence that the wall is bathed in 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 does not make it at all probable 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 much lower than the probability that the wall is red given only that it looks red. In this way, the background evidence about the lighting reduces the degree to which your sensory evidence supports the conclusion that the wall is red. This is one way to interpret Pollock's claim 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 for the conclusion 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 the conclusion? I'll argue that it's not, since the evidence that you're hypoxic makes no impact on your degree of evidential support for the conclusion. There is no sense in which this 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 by adding the premise that hypoxia caused you to botch the deduction

from premises to conclusion. This expanded set of premises continues to entail the conclusion that you have enough fuel.

Moreover, entailment is the strongest kind of evidential link between premises and conclusion. Arguments come in varying degrees of strength: the stronger an argument, the higher the probability that its conclusion is true given that all its premises are true.[5] 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*.[6]

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

What changes when you acquire the new evidence that you are hypoxic? You now 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, it 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 impair your eyesight or your ability to read the control panel. Rather, the point of the example is that hypoxia impairs your capacity to acquire new knowledge by deduction. Since the evidence of hypoxia doesn't change the evidential probability of the premises, and it doesn't undermine the entailment, it doesn't change the evidential probability of the conclusion.

I conclude that evidence of hypoxia is no evidential defeater at all, since it does nothing to reduce the evidential probability of your conclusion. All evidential defeaters, whether or not they provide evidence against a conclusion, reduce the degree to which your total evidence supports a conclusion. Given a probabilistic conception of evidential support, we can articulate this constraint as follows:

---

[5] Here, I'm following Skyrms (1966, Ch. 2), who proposes that the strength of an argument is measured by the conditional probability that its conclusion is true given that all its premises are true. This is what he calls "inductive probability".

[6] On a probabilistic conception of evidential support, probabilities are used to model degrees of evidential support. In this context, I have in mind an absolute rather than an incremental conception of evidential support.

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.[7]

The higher-order evidence that you are hypoxic doesn't satisfy this probabilistic constraint on evidential defeat, since it doesn't reduce the evidential probability that you have enough fuel.

Some epistemologists may respond to this objection by rejecting the probabilistic conception of evidential support on which it depends. In reply, however, the probability calculus is no more than a convenient framework for articulating the intuitive idea that entailment preserves strength of evidential support from premises to conclusion. Of course, someone might reject this claim too, but this comes at the cost of obscuring an important epistemic distinction between deduction and induction. Deductive entailment is the strongest kind of evidential link between premises and conclusion, since it preserves not only truth, but also degrees of evidential support.

I conclude that the higher-order evidence that you are hypoxic doesn't destroy your knowledge that you have enough fuel by defeating your first-order evidence for this conclusion. This is not to deny the intuitive datum that it destroys your knowledge in some other way. Moreover, I have no complaint about the common practice of articulating this datum using the language of 'defeat' just so long as we 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 *how* higher-order evidence of hypoxia destroys your knowledge that you have enough fuel. In this section, I've argued that it doesn't destroy your knowledge by reducing your 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*. On this view, higher-order evidence can be "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: that is, 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

---

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

evidence supports believing that you have enough fuel, while your higher-order evidence supports believing 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, p. 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.[8] 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*, according to which you should believe whatever is supported strongly enough by your total evidence.

Christensen's proposal can be understood as a form of *bifurcationism* about epistemic rationality.[9] On this view, the structural requirements of coherence, including meta-coherence, are distinct from and irreducible to the substantive 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, fc). 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:

---

[8] See Feldman and Conee (1985) for a classic statement of evidentialism.
[9] See Smithies (fc) for criticisms of bifurcationism, including Christensen's version, and a defense of unificationism.

Probabilistic Accessibilism: Necessarily, if the evidential probability that *p* is *n*, then it is evidentially certain that the evidential probability that *p* is *n*.[10]

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 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 probably 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 meta-incoherent beliefs of the form, '*p* and my evidence probably doesn't support *p*'. Hence, respecting your evidence guarantees meta-coherence.

On this view, 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

---

[10] See Smithies (2019, Chs. 7-11) in defense of probabilistic accessibilism. Christensen (2010b) and Elga (2013) reject probabilistic accessibilism in favor of rational reflection principles, but these higher-order constraints are too weak to prohibit forms 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.

made certain by your evidence. In general, 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 you have misleading higher-order evidence about your response to your evidence, your total 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 an analogy from the practical domain. 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, since he has evidence that his motivation is selfish. In fact, however, this evidence is misleading, and his action is motivated by altruistic moral reasons. This is not an example of akratic action. There need be no incoherence in doubting your own motivation for action. The philanthropist can coherently think to himself, "Perhaps I'm doing this for the wrong reasons, but there's not much I can do about that, and it's still the right thing to do."

The same point applies in the epistemic domain. Suppose a loving mother is convinced that her son is innocent of a crime for which he has been charged. She knows her belief is supported by good evidence, but she suspects her belief is based on bad reasons, since she has evidence that she is motivated by wishful thinking. In fact, however, this evidence is misleading, and her belief is rationally responsive to good evidence. This is not an example of akratic belief. There need be no incoherence in maintaining your beliefs while doubting their motivation so long as you know they are supported by good evidence. The mother can coherently think to herself, "Perhaps I believe this for the wrong reasons, but there's not much I can do about that, and it's still the right thing to believe."

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

(a)   It's certain that my evidence supports believing that I have enough fuel.
(b)   But my belief that I have enough fuel is probably based on bad reasoning.

(c)   So, I probably got lucky because bad reasoning led me to form a belief that is supported by the evidence.[11]

In the absence of strong evidence, it's irrational to believe in a lucky coincidence. After all, the prior probability of a coincidence is very low. Even so, coincidences do sometimes happen. And we sometimes have strong evidence that they happen. Given strong enough evidence, the posterior probability of a coincidence may be extremely high. There is nothing in principle to rule out the possibility that your evidence supports the conclusion that you got lucky by believing the right thing for the wrong reasons. As we've seen, these cases do happen. In my view, the hypoxia example is another case in point.

Does this mean dogmatism is a rational response to misleading higher-order evidence? I don't think so. Consider our steadfast pilot who maintains his belief that he has enough fuel in the face of the evidence that he is hypoxic. Suppose he concludes that he must have got lucky, since his belief is supported by evidence although it is based on bad reasoning. Intuitively, this merely compounds his irrationality. It was already irrational for him to retain his first-order belief that he has enough fuel, and this 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 adding irrationally dogmatic higher-order beliefs to his irrationally dogmatic first-order beliefs.

This just goes to show 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 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, the problem iterates 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 more specific reasons to doubt that the pilot's evidence supports meta-incoherent beliefs. We must 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; fc) argues that the steadfast pilot has *unreasonable knowledge*.[12] On this view, the pilot retains his knowledge that he has enough fuel

---

[11] Horowitz (2014, pp. 9-10) makes a related point, although her objection targets views on which your total evidence supports believing, "I probably got lucky because my true belief is based on misleading evidence."

[12] Lasonen-Aarnio (2010) focuses on standard cases of undercutting defeat, but she extends her proposal to higher-order evidence in Lasonen-Aarnio (fc).

when he acquires the misleading higher-order evidence that he is hypoxic. The only 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 the evidence 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, p. 2)

The steadfast pilot is unreasonable because he fails to manifest dispositions 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 in the "bad case" in which it is accurate. Since he is actually in 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. 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 now unreasonable.

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 these 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. It is counterintuitive to suppose that the steadfast pilot retains both his first-order knowledge that he has enough fuel and his higher-order knowledge that his evidence supports this conclusion. And it is even more counterintuitive that he

can use this knowledge, together with the misleading evidence that he is hypoxic, to rationally conclude that he probably got lucky on this occasion.

Lasonen-Aarnio explains away conflicting intuitions by appealing to the error theory that we tend to conflate reasonableness and knowledge. I agree that our intuitions in 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 infer 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 think 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 have knowledge, it's not enough that your belief is true: it must be *reliable* in the sense that it manifests a more general disposition to have true beliefs. Similarly, to be justified or rational, 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 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 expresses some sympathy for 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. And yet the case of mathematical knowledge suggests that this simple view needs refinement. I cannot know that Fermat's Last Theorem is true just by wishful thinking, but there is no danger that my belief is false, since its content is necessarily true. Arguably, this is because my safe belief has an unsafe basis: it is based on exercising an unreliable disposition that could easily yield false beliefs in other propositions. As Sosa (2003, pp. 137-40) puts the point, a belief is knowledge only if it is not only safe from error but also *virtuous* in the sense that it manifests a more generally reliable disposition to form safe beliefs.[13]

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. For example, knowledge is often thought to set normative standards for belief and action: if you know that p, then it is rationally permissible 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

---

[13] Elsewhere, I extend this point from knowledge to justified belief by considering the problem of the speckled hen (Smithies 2019, pp. 349-60).

enough fuel: his loved ones can be justly angry that he acted so recklessly. As Lasonen-Aarnio acknowledges: "Subjects who retain knowledge in defeat cases are genuinely criticisable" (2010, p. 15).

We cannot explain the pilot's culpability, while also maintaining that he retains knowledge, without abandoning a plausible theoretical connection between knowledge and blameworthiness. It is plausible that you can be blameless for your beliefs despite lacking knowledge, but it is not so plausible that you can be blameworthy for your beliefs despite having knowledge. This devalues knowledge in ways that should be unattractive to anyone and perhaps especially to proponents of knowledge-first epistemology.

To conclude, we need to explain how higher-order evidence can not only make your beliefs unreasonable but can also destroy your knowledge. Although Lasonen-Aarnio's proposal fails to explain this datum, it nevertheless contains an important insight about the importance of dispositions in epistemic evaluation.[14] This point figures prominently in the discussion to follow, since I argue that we need to invoke facts about the reliability of your doxastic dispositions in order to explain how higher-order evidence destroys knowledge.

4. The Proper Basing Relation

My own explanation of how higher-order evidence destroys first-order knowledge appeals to *the proper basing relation*.[15] 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. A belief is justified only if it is properly based on good evidence and only justified beliefs can be knowledge.

We can articulate the point in terms of the familiar distinction between propositional and doxastic justification.[16] 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. Hence, *proper basing* is the relation between a belief and a body of 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. Elsewhere, I put

---

[14] See Lasonen-Aarnio (fc) for her current views on dispositional evaluations.

[15] Compare Smithies (2015; 2019, Ch. 10). Similar views are defended by van Wietmarschen (2013) and Silva (2017), but I propose a distinctive explanation of how higher-order evidence undermines proper basing.

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

this by saying that this higher-order evidence is a *doxastic defeater*, rather than a *propositional defeater* (Smithies 2015, p. 2786).

When is a belief properly based on supporting evidence? It's not enough that my belief is based on evidence that happens to support the belief. My belief must 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.[17]

Dogmatic beliefs violate this condition. Consider my belief that drinking red wine in moderation 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 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 discover new studies casting doubt on the health benefits of drinking red wine. Moreover, this is something that could easily happen. Intuitively, I don't know that drinking wine is good for me, even if my belief is true and based on good evidence. This is because my belief is not *properly based* on my evidence, since it manifests a disposition that is not sensitive enough to changes in what my evidence supports.

I propose 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 evidence, because it is not *properly based* on his evidence. The 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 it is accurate because he really has made a mistake in reasoning.

We can set aside modally remote cases in which the steadfast pilot is actually suffering from hypoxia, since this dramatically changes his capacity for reasoning. Instead, let's consider more quotidian cases in which his reasoning capacities is held constant. Since his capacity for reasoning is fallible, there are bad cases in which he makes routine errors in calculation through a defective exercise of the very same capacities he employs in the good case. Let's consider a bad case in which he makes a routine error, and his co-pilot points out the mistake. Again, this is something that could easily happen. If he is disposed to remain steadfast in

---

[17] In other work, I use the problem of the speckled hen to motivate this constraint on proper basing (Smithies 2015; 2019, Ch. 11). More specifically, I propose 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.

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 spite 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 despite retaining his knowledge that he has enough fuel. In contrast, my goal here is to argue that his belief is doxastically unjustified, and so fails to be knowledge, since it is not properly based on supporting evidence. In the good case, his belief is true and supported by evidence, whereas in the bad case, it is false and unsupported by evidence. So, even in the good case, the steadfast pilot lacks knowledge, since his belief manifests a disposition that is insufficiently 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.[18]

We can apply the same reasoning one level up to explain why the steadfast pilot lacks *higher-order* knowledge about his own epistemic situation. Suppose he believes not just that he has enough fuel but also that his evidence supports this conclusion. Although this higher-order belief is true and entailed by his 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.

To be clear, I am not saying that that the steadfast pilot *loses* his knowledge when he acquires the higher-order evidence that he is hypoxic. As Lasonen-Aarnio (2010, pp. 3-8) explains, this claim is hard to sustain. How exactly is his knowledge destroyed when he acquires this higher-order evidence? There is no relevant change in what his evidence supports, since his new evidence continues to support the conclusion that he has enough fuel. Moreover, there is no relevant change in his response to his evidence, since he is steadfast enough that acquiring this new higher-order evidence makes no difference to the basis on which his belief is held. So how can this new higher-order evidence destroy his knowledge?

My response is that the steadfast pilot doesn't *lose* his 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 new higher-order evidence that he is hypoxic because he responds by abandoning his belief that he has enough fuel. There is no knowledge without belief.

The key point is that knowledge sometimes requires being disposed to abandon belief in response to new evidence. The steadfast pilot doesn't have

---

[18] To undermine knowledge, there has to be a close enough case in which the agent ignores accurate higher-order evidence that she has made a mistake, but I doubt we can specify in non-epistemic terms what counts as a close enough case.

knowledge because he is disposed to retain his belief in the face of higher-order evidence whether it is misleading or accurate. In contrast, the conciliatory pilot has knowledge only because he is disposed to respond to such higher-order evidence by abandoning his belief. The conciliatory pilot is more reliable than the steadfast pilot in responding to evidence 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 for anyone to retain knowledge in the hypoxia case, but only that it's impossible in practice for creatures like us. I see no reason to rule out the metaphysical possibility of ideally rational agents who have misleading evidence that they are suffering from hypoxia. Because ideally rational agents are perfectly sensitive to what their evidence supports, they can remain steadfast in good cases without thereby manifesting any disposition to remain steadfast in bad cases where their reasoning dispositions are held constant. And yet this is beyond our limited human capacities, since we are only imperfectly sensitive to our evidence. Any human who remains steadfast in the good case thereby manifests some disposition to remain steadfast in the bad case too.

We can now draw a more general conclusion about the relationship between propositional and doxastic justification. There are cases where your evidence gives you propositional justification to believe a conclusion, although you are unable to form 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.[19]

There are independent reasons to reject the doxastic constraint. Suppose you ingest a reason-distorting drug that temporarily impairs your ability to form beliefs and other doxastic attitudes that are properly based on your evidence. Whatever doxastic attitudes you adopt, they are guaranteed to be doxastically unjustified. It is implausible that 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 by default you have justification to remain agnostic. Even so, there is no guarantee that you can always form a doxastic attitude that is properly based on what your evidence supports. After all, you might 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 propositional justification into doxastic justification. If there are doxastic constraints on propositional justification,

---

[19] Proponents include Goldman (1979) and Turri (2010).

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 square 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 supports. It is more plausible that ingesting the drug compromises your epistemic rationality by disabling you from responding appropriately to your evidence and thereby converting propositional justification into doxastic justification.

On a plausible version of evidentialism, there are no doxastic constraints on propositional justification. What your evidence supports is one question, but it's another question whether you can form beliefs that are supported by and properly based on your evidence. One theoretical cost of conflating these questions is to obscure the epistemic function of higher-order evidence.

5. Ideal and 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.

The negative datum was explained 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 rather puzzling in 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.[20] By ideal standards, epistemic rationality always requires respecting your evidence. This ideal sometimes outstrips our limited human capacities: we are not always capable of responding rationally to our evidence. Moreover, we know this; or we should know this, since we all have compelling evidence of our own cognitive limitations. By non-ideal standards, in contrast, epistemic rationality sometimes permits (indeed, requires) responding to such higher-order evidence by adopting policies that diverge from the 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.

---

[20] Compare Smithies (2015; 2019, Ch. 10).

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 change which conclusions are supported by your total evidence. Rather, it determines which response to your total evidence is required by non-ideal standards of epistemic rationality.

This is an instance of a more general point. It doesn't always make sense to try to achieve what you know would be best if only you were to succeed. This is because you sometimes know – or have good evidence – that your attempt to achieve the optimal outcome may fail. In such cases, trying to do the best thing sometimes 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.[21] 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 the follow the rule doesn't guarantee that you will succeed. When you have evidence that you may fail in your attempt to follow a rule, the expected consequences of trying to follow the rule can diverge from the expected consequences of successfully following the rule. In such cases, the best rule to follow is not always the best rule to try to follow.

Consider the case of Professor Procrastinate (Jackson and Pargetter, 1986). When is invited to review a book, he has three options available to him, which are listed below in rank order from best to worst:

R1. Accept the invitation and complete the review on time.
R2. Decline the invitation.
R3. 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 notoriously context-sensitive. Intuitively, 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 there is also a sense in which he should decline the invitation, since he knows (or perhaps has misleading evidence) that he is unlikely to complete the review on time: if he tries to achieve the best outcome, then he is more likely to bring about the worst

---

[21] Here, I follow Lasonen-Aarnio (2010, pp. 14-15) and Schoenfield (2015, pp. 650-3), who both use rule consequentialism to draw normative distinctions, but see below for some important strategic differences.

outcome. Instead of 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. R1 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 R1 is that he will end up following R3 instead. The expected consequence of trying to follow R2, in contrast, is that he will succeed. And there is greater expected value in following R2 than following R3. Hence, the expected value of trying to follow R2 is greater than R1, although the expected value of successfully following R1 is greater than R2. This is why ideal standards require following R1, although non-ideal standards require following R2.

We can extend the same framework to *epistemic rules*: that is, rules for forming and revising beliefs. In epistemic evaluation, we're concerned solely with the *epistemic value* of the expected consequences of following (or trying to follow) a rule. And when we evaluate rules for *epistemic rationality*, we're concerned with a specific dimension of epistemic value: namely, how well you succeed in proportioning your beliefs to your evidence. Finally, we're concerned only with the *direct* consequences that are expected to result from following (or trying a follow) an epistemic rule. As Sophie Horowitz (2019, p. 116) observes, it's not epistemically rational for an agent with severe arachnophobia to follow the rule, "Never believe there is a spider nearby," just because this indirect strategy helps him to remain calm and follow other epistemically rational rules.

By ideal standards, epistemic rationality requires following the evidentialist rule: "Always proportion your beliefs to your evidence!" Assuming evidentialism, this is the best rule to follow, since the expected outcome of following the rule is perfect epistemic rationality. 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 we have evidence that we are likely to fail. There is sometimes greater expected epistemic value in trying to follow an alternative rule when you have evidence that adopting this strategy increases your responsiveness to evidence. In such cases, trying to follow the evidentialist rule is a kind of epistemic self-sabotage: it's a counterproductive strategy for maximizing your responsiveness to evidence. This is why non-ideal standards of epistemic rationality sometimes prohibit 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:

> *Steadfastness*: Maintain your first-order belief that you have enough fuel and your higher-order belief that your evidence supports this conclusion.

> *Level Splitting*: Maintain your first-order belief that you have enough fuel, while abandoning your higher-order belief that your evidence supports this conclusion.

Conciliation: Abandon your first-order belief that you have enough fuel and your higher-order belief that your evidence supports this conclusion.

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

It's clear enough that the *expected pragmatic value* of trying to follow Conciliation in the hypoxia case is greater than the alternatives. After all, the expected outcome of trying to follow either Steadfastness or Level Splitting is that you run the risk of an early demise, which is not worth taking for a scenic detour. But why think there is any discrepancy in *expected epistemic value*?

Miriam Schoenfield (2015, p. 652) and Maria Lasonen-Aarnio (fc, §5) argue that trying to follow Conciliation has greater expected epistemic value than the alternatives because they run the risk of false belief in bad cases where your higher-order evidence is accurate. But this strategy faces a problem. Why isn't the risk of false belief in bad cases neutralized by the chance of knowledge or true belief in good cases where your higher-order evidence is accurate? And why is this neutral outcome any worse than a consistent policy of agnosticism? Of course, we might assume that the value of knowledge or true belief in good cases is outweighed by the disvalue of false belief in bad cases, but I see no compelling basis for this assumption. My own strategy is different: it appeals to the disvalue of unjustified belief in good cases, rather than false belief in bad cases.[22]

The expected consequence of trying to follow Steadfastness in the hypoxia case is that you are like the *steadfast pilot*. As we saw in §4, the steadfast pilot cannot know or rationally believe that he has enough fuel. This is because his belief is held dogmatically in a way that is insensitive 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 are like the *conciliatory pilot*. The conciliatory pilot rationally believes and knows that he has enough fuel so long as he is disposed to become agnostic when he acquires the higher-order evidence that he is hypoxic. Hence, there is greater

---

[22] Lasonen-Aarnio (2010, fc) disagrees, of course, since she maintains that the steadfast pilot has knowledge in the good case. Schoenfield's (2015) strategy is different too, since she evaluates epistemic rules in terms of expected accuracy, rather than expected epistemic rationality.

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 what your evidence supports.

What about Level Splitting? The expected consequence of trying to follow this rule is that your first-order belief is like the steadfast pilot's, while your higher-order belief is like the conciliatory pilot's. Like the conciliatory pilot, you have higher-order knowledge that your evidence supports the conclusion that you have enough fuel so long as you are disposed to become agnostic when you acquire the evidence that you are hypoxic. Like the steadfast pilot, however, you lack the first-order knowledge that you have enough fuel, since your first-order belief is held dogmatically in a way that is insensitive to changes in your evidence.

As we have seen, trying to follow Conciliation has greater expected value than the alternatives because it yields knowledge and justified belief in the good case so long as you are disposed to become agnostic when you acquire the evidence of hypoxia. This is why you are required by non-ideal standards of epistemic rationality to become agnostic in response to this higher-order evidence, rather than believing what your evidence supports. Thus, we can explain the positive datum as well as the negative datum.

In summary, the epistemic function of higher-order evidence is twofold. First, it prevents you from responding properly to your evidence in the way that is required by ideal standards of epistemic rationality. And second, it affects which response to your evidence is required by non-ideal standards of epistemic rationality that are sensitive to evidence about your cognitive limitations.

6. Conclusions

Why prefer my account of the epistemic function of higher-order evidence to the more familiar alternatives discussed in §§1-3?

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. The grain of truth in this view is that the steadfast pilot manifests bad dispositions when he retains his belief in the face of misleading higher-order evidence. But the problem is that manifesting these bad dispositions is incompatible 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 it's 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. The grain of truth in this view is that your knowledge is destroyed when you acquire the higher-order evidence that you are hypoxic. The problem is that this is not plausibly 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. The grain of truth in this view is that epistemic rationality somehow requires "bracketing" your first-order evidence when you have misleading higher-order evidence. The problem is that this is not 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.[23]

References
Christensen, D. (2007.) Does Murphy's Law apply in epistemology? Self-doubt and rational ideals. *Oxford Studies in Epistemology*, 2, 3-31.

Christensen, D. (2010a.) Higher-order evidence. *Philosophy and Phenomenological Research*, 81 (1), 185-215.

Christensen, D. (2010b.) Rational reflection. *Philosophical Perspectives* 24: 121-140.

Elga, A. (2008.) Lucky to be rational. *Bellingham Summer Philosophy Conference*, unpublished manuscript.

Elga, A. (2013.) The puzzle of the unmarked clock and the new rational reflection principle. *Philosophical Studies*, 164 (1), 127-139.

Feldman, R. (2005.) Respecting the evidence. *Philosophical Perspectives*, 19, 95-119.

Feldman, R., & Conee, E. (1985.) Evidentialism. *Philosophical Studies*, 48 (1), 15-34.

Firth, R. (1978.) Are epistemic concepts reducible to ethical concepts? In A. Goldman & J. Kim (Eds.), *Values and morals* (pp. 215-229). Dordrecht: Kluwer.

Goldman, A. (1979.) What is justified belief? In G. Pappas (Ed.), *Justification and knowledge* (pp. 1-25). Dordrecht: Reidel.

Hawthorne, J., & Stanley, J. (2008.) Knowledge and action. *Journal of Philosophy*, 105 (10), 571-590.

Horowitz, S. (2014.) Epistemic akrasia. *Noûs*, 48 (4), 718-744.

Horowitz, S. (2019.) Predictably misleading evidence. In M. Skipper (Ed.), *Higher-order evidence: new essays* (pp. 105-123). New York, NY: Oxford University Press.

Jackson, F., and Pargetter, R. (1986.) Oughts, options, and actualism. *Philosophical Review*, 95 (2), 233-255.

Jackson, L. (2019.) How belief-credence dualism explains away pragmatic encroachment. *Philosophical Quarterly* 69: 511-533.

Kotzen, M. (2019.) A formal account of evidential defeat. In B. Fitelsen, R. Borges, & C. Braden (Eds.), *Themes from Klein: knowledge, justification, and skepticism* (pp. 213-234). Cham: Synthese Library.

Lasonen-Aarnio, M. (2010.) Unreasonable knowledge. *Philosophical Perspectives*, 24, 1-21.

Lasonen-Aarnio, M. (Forthcoming.) Dispositional evaluations and defeat. In J. Brown & M. Simion (Eds.), *Reasons, justification, and defeat*. New York, NY: Oxford University Press.

Pollock, J. (1986.) *Contemporary theories of knowledge*. London: Hutchinson.

Schoenfield, M. (2015.) Bridging rationality and accuracy. *Journal of Philosophy*, 112 (12), 633-657.

Silva, P. (2017.) How doxastic justification helps us solve the puzzle of misleading higher-order evidence. *Pacific Philosophical Quarterly* 98: 308-328.

Skyrms, B. (1966.) *Choice and chance: an introduction to inductive logic*. Belmont, CA: Dickenson.

Smithies, D. (2015.) Ideal rationality and logical omniscience. *Synthese*, 192 (9), 2769-2793.

Smithies, D. (2019.) *The epistemic role of consciousness*. New York, NY: Oxford University Press.

Smithies, D. (Forthcoming.) The unity of evidence and coherence. In *Epistemic dilemmas*, edited by Nick Hughes. New York, NY: Oxford University Press.

Sosa, E. (2003.) Beyond internal foundations to external virtues. In L. BonJour & E. Sosa, *Epistemic justification: internalism vs. externalism, foundations vs. virtues*. Oxford: Blackwell.

Turri, J. (2010.) On the relationship between propositional and doxastic justification. *Philosophy and Phenomenological Research*, 80 (2), 312-326.

van Wietmarschen, H. (2013.) Peer disagreement, evidence, and well-foundedness. *Philosophical Review*, 122.3, 395-425.