

Respecting *All* The Evidence

Paulina Sliwa
University of Cambridge
Sidgwick Avenue
Cambridge, UK CB3 9DA
pas70@cam.ac.uk

Sophie Horowitz
Rice University
6100 Main MS-14
Houston, Texas 77005
sophie.horowitz@rice.edu

Respecting *All* The Evidence

Abstract: Plausibly, you should believe what your total evidence supports. But cases of misleading higher-order evidence -- evidence about what your evidence supports -- present a challenge to this thought. In such cases, taking both first-order and higher-order evidence at face value leads to a seemingly irrational incoherence between one's first-order and higher-order attitudes: you will believe P, but also believe that your evidence doesn't support P. To avoid sanctioning tension between epistemic levels, some authors have abandoned the thought that both first-order and higher-order evidence have rational bearing. This sacrifice is both costly and unnecessary.

We propose a principle, Evidential Calibration, which requires rational agents to accommodate first-order evidence correctly, while allowing rational uncertainty about what to believe. At the same time, it rules out irrational tensions between epistemic levels. We show that while there are serious problems for some views on which we can rationally believe, "P, but my evidence doesn't support P", Evidential Calibration avoids these problems. An important upshot of our discussion is a new way to think about the relationship between epistemic levels: why first-order and higher-order attitudes should generally be aligned, and why it is sometimes -- though not always -- problematic when they diverge.

Keywords: Higher-order evidence, rationality, reliability, epistemic akrasia, epistemic levels

1. Three desiderata for an account of higher-order evidence

Consider the following story:

Calculation: Anton is an anesthesiologist, trying to determine which dosage of pain medication is best for his patient: dose A, B, C, or D. To figure this out, Anton assesses some fairly complex medical evidence. When evaluated correctly, this kind of evidence determines which dose is right for the patient. After thinking hard about the evidence, Anton becomes highly confident that dose B is right. In fact, Anton has reasoned correctly; his evidence strongly supports that B is the correct dose.

Then Sam, the chef at the hospital's cafeteria, rushes in. "Don't administer that drug just yet," he says guiltily. "You're not in a position to properly assess that medical evidence. I slipped some reason-distorting mushrooms into your frittata earlier as a prank. These mushrooms make you much less reliable at determining which dose the evidence supports: in the circumstances you presently face -- evaluating this type of medical evidence, under the influence of my mushrooms -- doctors like you only tend to prescribe the right dose 40% of the time!" In fact, Sam is mistaken: the mushrooms he used were just regular dried porcini, and Anton's reasoning is not impaired in the least. But neither he nor Anton knows (nor has reason to suspect) this.

Suppose Anton justifiably believes that Sam has given him the mushrooms. How should Anton accommodate his evidence about which dose of medication to give to his patient? And how should he accommodate Sam's testimony?

Plausibly, Anton should believe what his evidence supports. Recent work in epistemology commonly divides evidence like Anton's into two categories: "first-order" and

“higher-order”. Anton’s first-order evidence, which bears directly on the question of which dose is right for his patient, consists of facts about the patient: his weight, age, etc. His higher-order evidence -- Sam’s testimony -- bears directly on the reliability of Anton’s reasoning. Intuitively, Anton’s final doxastic state should reflect both kinds of evidence. This suggests the following two desiderata for a plausible account of what Anton should believe:

Desideratum 1: One’s rational beliefs should reflect the bearing of one’s first-order evidence.

Desideratum 2: One’s rational beliefs should reflect the bearing of one’s higher-order evidence.

These desiderata look straightforward enough. But they aren’t obviously compatible with one another. On the one hand, it seems like in order to accommodate his first-order evidence, Anton should believe *that dose B is right*. On the other hand, Anton’s higher-order evidence -- Sam’s testimony -- suggests that doctors in Anton’s situation are only 40% likely to pick the right dose. So it seems that in accommodating his higher-order evidence, Anton should become quite uncertain as to whether *his evidence supports dose B*. If Anton were to take both kinds of evidence at face value, his first-order and higher-order attitudes would exhibit an odd kind of incoherence. He might believe something like this: “B is the right dose. But I’m not sure what my evidence supports. Maybe it supports dose B, but maybe not.”

It is implausible that believing Moore-paradoxical combinations like this one could ever be rational.¹ This observation suggests a third desideratum for a plausible account of what someone in Anton’s position should believe:

Desideratum 3: One’s rational first-order and higher-order doxastic attitudes should not be in tension.

According to Desideratum 3, the right way to accommodate first-order and higher-order evidence should not lead you to believing a Moore-paradoxical conjunction.

Evidential situations like Anton’s, in which an agent’s first-order evidence and higher-order evidence come apart, pose a challenge for accounts of higher-order evidence. Each of our three desiderata is initially plausible, but it seems impossible to systematically and straightforwardly satisfy all of them. Many recent accounts have opted to give up one or

¹See Smithies [2012] for further discussion of this point.

another completely, and much of the recent literature could be described as a debate over *which* desideratum should be abandoned.

The purpose of this paper is to motivate and explore a view of the rational effect of higher-order evidence that respects all three desiderata. Our discussion will proceed as follows. In the first half, we survey some recent accounts of higher-order evidence that have abandoned one of our three desiderata in order to accommodate the other two. We argue that this comes at a serious cost. In the second half, we propose and defend a principle – Evidential Calibration – describing what agents in situations like Anton’s should believe. If Evidential Calibration is right, it turns out that tension between one’s first- and higher-order attitudes can sometimes be rational. But, we argue, Evidential Calibration still has the resources to explain the main motivation behind Desideratum 3: that it’s irrational to believe “P, but my evidence doesn’t support P.” An advocate of Evidential Calibration can distinguish between rational and irrational cases of tension between epistemic levels, while maintaining that there are genuine instances of both.

2. Motivating the three desiderata

In the last section we introduced three desiderata for an account of the effect of higher-order evidence. In this section, we will motivate these desiderata by looking at problems for recent views that have given up one desideratum entirely in order to accommodate the other two.

2.1 Denying the import of higher-order evidence

Let’s return to Calculation. Here is one reaction you might have to Anton’s evidential situation: Since Anton’s first-order evidence supports dose B, he should believe that dose B is right. And since Anton knows that dose B is right, and that he came up with this answer after looking at his evidence, he should also be completely certain that he has accommodated his evidence correctly.

This view is a natural extension of Kelly [2005]’s account of peer disagreement.

Kelly writes:

“[O]nce I have thoroughly scrutinized the available evidence and arguments that bear on some question, the mere fact that an epistemic peer strongly disagrees with me about how that question should be answered does not itself tend to undermine the rationality of my continuing to believe as I do. ... Indeed, confidently retaining my

original belief might very well be the uniquely reasonable response in such circumstances.”²

One might motivate this response as follows: if you have first-order evidence relevant to a particular question, your higher-order evidence will be trumped or screened off by this first-order evidence. What your evidence supports is an a priori matter, so if you have accommodated your first-order evidence correctly, you should be certain of this. In the presence of higher-order evidence, then, a rational agent should hold on to her initial first-order attitude. She should also maintain her high confidence that her evidence does support the first-order attitude in question, and that this first-order attitude is rational.

On this view, Anton’s first-order evidence clearly has rational import. So this view respects Desideratum 1. It also respects Desideratum 3: since Anton should be certain that his first-order attitude is rational, his first-order and higher-order attitudes will not be in tension. (Rather than the Moore-paradoxical “P, but by evidence doesn’t support P”, Anton should believe, “P, and my evidence supports it.”) But this view does not respect Desideratum 2: Anton’s higher-order evidence has no rational bearing on what he should believe. If Anton is rational, on this view, he should ignore Sam’s testimony completely.

Giving up Desideratum 2 incurs some serious costs. If the view under consideration is right, then after hearing Sam’s testimony, it would be rational for Anton to respond like this: “Sure, Sam, you’ve drugged me; but I must be immune to the drugs’ effects. After all, I’m highly confident that dose B is right, *and* my evidence supports it!” As many participants in the debate have pointed out, this seems irrational. If Anton takes Sam to be reliable, and believes that he has eaten the mushrooms, he should think it’s likely that his medical reasoning is impaired. Furthermore, he should think that he has probably reached a conclusion that is *not* supported by his evidence: if the mushrooms make doctors less reliable at prescribing the right dose, presumably they do so by impairing one’s ability to figure out which dose is right. So if Anton takes Sam’s testimony seriously, he should not be certain that his medical evidence supports what he thinks it supports.

2.2 Denying the import of first-order evidence

² Kelly [2005], p. 4-5. It is not entirely clear whether Kelly’s position here is better described as giving up Desideratum 2, or as giving up Desideratum 3 like the “Level-Splitting” views discussed below. Nevertheless, many of the considerations Kelly raises could naturally be used to motivate giving up Desideratum 2, so it is helpful to look at his view in this context. Also see Titelbaum [ms].

Christensen, Elga, Feldman, and Kelly (more recently) have all defended the position that higher-order evidence is epistemically significant: that in cases like Anton's, one is rationally required to reduce confidence in one's first-order attitudes in response to higher-order evidence. So after hearing Sam's testimony, Anton should become less confident that his first-order evidence supports dose B *and* less confident that dose B is right.

Thinking about why Anton should reduce confidence in dose B might lead one to the view that higher-order evidence does *all* the work in determining what Anton should believe. Consider the following line of thought: upon learning about the mushrooms, Anton should become less confident that B is the right dose. How much should Anton reduce confidence? We said that Sam told Anton that doctors in his position only prescribed the right dose 40% of the time. So a natural view to take is that after hearing Sam's testimony, Anton should only be .4 confident that dose B is right. Moreover, you might think, Anton should be certain that his *total* evidence supports .4 confidence in B. Once he has reduced his confidence that B is the correct dose, Anton should be certain that he has accommodated all of his evidence rationally.

White [2009] discusses a view along these lines, and suggests (though stops short of defending) the following as a general principle for how to rationally accommodate higher-order evidence:

Guess Calibration: If I draw the conclusion that P on the basis of any evidence E, my credence in P should equal my prior expected reliability with respect to P.³

Your *reliability* with respect to P is how likely you are to come to a true guess about P. White characterizes reliability like this:

"I suggest we think about [reliability] along the lines that we think about objective *chance*. ... This notion of reliability is a modal one, being logically independent of any actual performance. A reliable thermometer has an objective propensity to give accurate readings, even if by some unfortunate fluke it rarely does (conversely a hopelessly unreliable thermometer might by chance give accurate readings)."⁴

Guess Calibration deals with one's *prior expected reliability* in order to accommodate cases when you are uncertain about how reliable a certain process is. In those situations, you might consider a number of possibilities for how reliable the process is, and assign some credence

³White [2009]. White calls this the "Calibration Rule". White is not explicit about what he means by "drawing the conclusion"; sometimes he writes as if it amounts to forming a belief, but the thought behind the rule is general enough to also apply to cases where we should suspend judgment. We will interpret it as a view about educated guesses, so as to stay neutral about this question. (We will say more about how we understand "educated guesses" in Section 3, when we spell out our view, Evidential Calibration.)

⁴White [2009], p. 234.

to each. In those cases, you can find the *expected* reliability of the process by taking an average of all of your candidate levels of reliability, weighted by how likely you take each one to be right.

White points out that plausibly, for any process, your credence that the process will produce an accurate outcome should equal the expected reliability of that process. It's also plausible that, after seeing which outcome a process produced, your credence that this outcome is accurate should equal your prior expected reliability of the process given that this is the outcome. Guess Calibration simply extends this general line of thought to our own cognitive processes. (White writes: "Given my cognitive capacities, what evidence I have to go on, and various environmental factors, I have a certain propensity to form a true belief as to whether p."⁵) And if we treat our own reasoning as we treat any other indicator of the truth, its expected reliability should rationally affect how much we trust its deliverances.

According to Guess Calibration, higher-order evidence does have a rational effect on what one should believe in evidential situations like Anton's. (Since it lowers Anton's expected reliability, it also lowers his rational credence in the dose he guessed to be correct.) Guess Calibration therefore meets Desideratum 2. It also meets Desideratum 3. According to Guess Calibration, what Anton should believe is determined by two factors: the fact that he has drawn the conclusion that B is the correct dose, and Sam's testimony concerning his expected reliability. Plausibly, Anton could be rationally certain about what those two factors are, and about how Guess Calibration says they should be accommodated. So if Guess Calibration is right, Anton can be rationally certain that he has responded to his total evidence rationally.

But note that the fact that Anton's first-order evidence supports dose B, rather than some other dose, has dropped out of the picture; it has no rational bearing on what Anton should believe. Guess Calibration gives up Desideratum 1 completely. This has some troubling consequences. Suppose that, as before, Anton rationally concludes that B is the correct dose for his patient. In the next room over, Anton's colleague Anna evaluates the same first-order evidence – but unlike Anton, Anna irrationally concludes that dose A is right. Both Anton and Anna receive the same testimony from Sam concerning their unreliability. Then Anton and Anna revise their first-order attitudes: Anton becomes .4

⁵ Ibid.

confident that B is the correct dose, while Anna becomes .4 confident that A is the correct dose.

There is an important asymmetry between Anton and Anna: Anton accommodated their shared first-order evidence correctly, while Anna did not. But according to Guess Calibration, both Anton and Anna are fully rational: both drew a conclusion about which dose was correct, and both reduced confidence in that guess in response to Sam's testimony. So Guess Calibration cannot explain the asymmetry between the two.⁶

This problem becomes more pressing if we consider a very simple case. Imagine someone who has absolutely no reason to doubt her reliability at assessing her evidence. Perhaps she has never rationally erred before, and has no reason to think that she faces any special problems today (for instance, due to mushrooms). She may not even have a track record, or may not be aware of her track record. Free from these higher-order complications, today she is evaluating some evidence which supports P to degree .8.⁷ What should her credence in P be? Well, since she fully expects to accommodate this evidence rationally, her expected reliability in guessing truly about P should be .8 as well. *But* suppose that on this occasion, against the odds and for the very first time, she misjudges her evidence. She concludes \sim P rather than P, and ends up with .8 confidence in \sim P.

Is this response rational? Guess Calibration says that it is. But surely, the fact that an agent has no reason to doubt her reliability does not make *any* conclusion she comes up with rational; she is still rationally required to believe what her evidence supports. To deny that would be a drastic departure from normal ways of thinking about evidence. (That is, if the evidence supports P to degree .8, it's hard to deny that absent any other factors, one should have .8 confidence in P.) The right account of higher-order evidence should at least maintain that first-order evidence has its usual bearing in cases where there is no reason for higher-order doubt.

⁶ This is one of the main lines of criticism that has been raised in the disagreement and higher-order evidence literature against views on which higher-order evidence mandates reducing confidence in one's conclusions. See, e.g., Kelly [2010] for an example of this objection in the context of peer disagreement; see Christensen [2007 b] and [2011] for responses. See also Weatherson [ms] for more extensive criticism of views on which accommodating higher-order evidence requires us to "screen off" first-order evidence completely. Christensen ([2011], section 1) denies his "Conciliatory" view is committed to rejecting Desideratum 1. According to Christensen, "Conciliationism tells us what the proper response is to one particular kind of evidence. ... If one starts out by botching things epistemically, and then takes correct account of one bit of evidence, it's unlikely that one will end up with *fully rational beliefs*." (p. 4, highlighting ours) Our project can be seen as complementary: our aim is to articulate how one should respond to one's *total* evidence.

⁷ Here we are thinking of the simplest case: someone who has *no reason* to doubt her reliability. This may be distinct from someone who has *good reason* to think she's reliable. See Christensen [2011] and Vavova [ms] for more on this distinction.

So the right account of higher-order evidence should respect Desideratum 1: what's rational to believe should be determined – at least in part – by what one's first-order evidence actually supports.

2.3 Allowing tension between first-order and higher-order attitudes

We have seen that giving up on either Desideratum 1 or Desideratum 2 raises serious problems. Surveying this state of affairs, one might naturally wonder if there is a third way to go: perhaps the rational response in situations like Anton's is to have one's final doxastic state determined by *both* one's first-order *and* one's higher-order evidence. So, one might defend the following verdict in Calculation:

Anton is rationally required to be highly confident in dose B. Anton is also rationally required to be only .4 confident that his evidence supports that dose.

Views defending this verdict have become popular in recent literature.⁸ Weatherson, for instance, writes that an agent in a situation like Anton's "...should believe p, and she should believe that she has most reason to not believe p."⁹ According to this view, Anton's rational first-order attitude is solely determined by what his first-order evidence supports, and his rational higher-order attitude is solely determined by what his higher-order evidence supports. The view therefore respects Desiderata 1 and 2. But it gives up on Desideratum 3: Anton's first-order and higher-order attitudes are obviously in tension. For the remainder of the paper, we will refer to views that violate Desideratum 3 in this way as "Level-Splitting" views. According to Level-Splitting, in cases like Calculation, it can be rational to believe something of the form, "P, but my evidence doesn't support P", or "P, but I shouldn't believe P".

The Level-Splitting conclusion is intuitively odd and has seemed highly implausible to many authors: it is commonly held that doxastic states like this can never be rational. Some have taken this point to be so obvious, and denying it to be so beyond the pale, that they

⁸See Coates [2012], Hazlett [2012], Weatherson [ms], Lasonen-Aarnio [forthcoming], and Wedgwood [2011] for views roughly along these lines. Williamson ([notes a] and [notes b]) defends the view that one can know that P, but that it can be highly improbable on one's evidence that one knows that P. Coates and Weatherson focus on the rationality of believing conjunctions like, "P, but my evidence doesn't support P"; others accept verdicts along these lines as a consequence of their other commitments. We will not go through the details of what these authors say here, so what we say here should not be taken as definitive arguments against their particular views. Rather, we focus on the implications for a certain natural way of spelling out a Level-Splitting view, and (eventually) argue why we should not go this way.

⁹Weatherson [ms], p. 15. He continues: "And, assuming a natural connection between evidence and reason, that in turn isn't very different from it being the case that she should believe p, and she should believe that her evidence does not support p."

presuppose it or take it as a premise rather than offering an argument for it.¹⁰ But defenders of Level-Splitting views have argued that perhaps giving up Desideratum 3 isn't so terrible after all. They point out, first of all, that first-order evidence and higher-order evidence directly bear on very different subject matters. Anton's first-order evidence bears directly on anesthesia doses, while his higher-order evidence bears directly on Anton's rationality. Why think that there should be a general rational constraint on how one's attitudes about these two completely different subject matters should relate? Secondly, similar kinds of level-tension are already familiar from ethics. There's nothing odd about endorsing conjunctions such as "Sally ought to ϕ , but she ought to believe that she has most reason not to ϕ ": so epistemic versions, some argue, should not seem so odd either.¹¹

One might also find a Level-Splitting view plausible because of skepticism about Desideratum 3 itself. Full compliance with Desideratum 3 seems to rule out any rational doubt about whether one's doxastic states are rational. But plausibly we can be rationally uncertain about what our evidence supports. Sometimes we may even be rationally uncertain about what our evidence *is*.¹² Calculation looks like a paradigmatic case of such rational uncertainty. So, one might think, insofar as Desideratum 3 is incompatible with rational uncertainty in Calculation, we should just give up on it altogether.

Whatever considerations tell in favor of Level-Splitting, though, the intuitive costs of the view are significant. (We will discuss this issue in more depth in section 4.) We can see this especially clearly when we consider how Anton should act in light of his doxastic state. Imagine Anton responding like this: "Sure, Sam, you've drugged me; it's probably irrational for me to be so confident that dose B is right for my patient. But B is right, so who cares if I'm rational? Fetch me that syringe!"¹³ This response seems patently irrational. Normally, if we are striving to do what's right, we act in the way we think is best -- and normally that corresponds to what our evidence supports. So surely Anton shouldn't be so confident that he's doing the right thing. Maybe instead he should feel guilty as he reaches for the needle: "I shouldn't be doing this", he might say; "I have no right to be so confident that this dose is correct. But it *is* correct, so here I go..." As much as it seemed irrational for Anton to cheerfully administer the medicine, doing so reluctantly is just as bad: if Anton is so

¹⁰See, e.g., Smithies [2012] and Titelbaum [ms]. Feldman [2005] defends a view broadly in line with Desideratum 3, and writes that "[o]ne wonders what circumstances could make [level-tension] reasonable" (p. 108-9).

¹¹ See Coates [2012] and Weatherson [ms] for this point, and for analogies to ethics.

¹² See Williamson [2000], especially Chapter 4.

¹³ This case parallels Christensen's "Reasonable Prudence" (Christensen [2010], p. 12).

confident that dose B is right, he should think that giving dose B *is* the right thing to do. What is the most rational way for Anton to react, then? Either way, treating the first-order and higher-order propositions as completely different subject matters leads to strange results.

3. Evidential Calibration

So far we have argued that views that give up on one of the three desiderata incur intuitive costs. We will now introduce a new principle and show how it accommodates Desiderata 1 and 2; in the next section we will discuss how it accommodates Desideratum 3.

We suggest that after hearing Sam's testimony, Anton should be .4 confident that dose B is right. He should be .4 confident because he expects to be 40% reliable; he should be .4 confident in *dose B* because that's what his first-order evidence actually supports. More specifically, we want to consider the following principle as a constraint on an ideally rational agent's credences:

Evidential Calibration: When one's evidence favors P over the alternatives to P, one's credence in P should equal the expected reliability of one's educated guess that P.

A rational agent following Evidential Calibration will match her credence in the proposition favored by her evidence to her expected reliability in making an educated guess that P. Together with the natural thought behind Desideratum 2 – that higher-order evidence should sometimes rationally affect our beliefs – Evidential Calibration tells us how we should revise our confidence in first-order propositions on the basis of higher-order evidence.

You may recognize Evidential Calibration as a close relative of White's Guess Calibration. Indeed, Evidential Calibration captures many of its cousin's benefits. But, crucially, Evidential Calibration *also* respects Desideratum 1. This is because Evidential Calibration gives a different explanation for why Anton should be .4 confident in *dose B*. According to Evidential Calibration, Anton should be .4 confident in B because *it's what his first-order evidence supports* and not – as Guess Calibration had it – because B is the conclusion he came up with.

The details of Evidential Calibration will require some unpacking. First of all, what is it for one's evidence to *favor* P with respect to its alternatives? By this we mean that the rational credence in P is at least as high as the rational credence in any of the alternatives to P, where P and its alternatives exhaust all the logical possibilities (that is, they form a partition). Suppose you are a meteorologist, and you are looking at some evidence to determine whether it will rain. You know that either Rain or \sim Rain must be true. Your

evidence *favors Rain* if it rationalizes a higher credence in Rain than in \sim Rain. Between two options, favoring P will just be a matter of rationalizing greater than .5 credence in P. Now suppose that you want to make a more fine-grained weather prediction about which of Snow, Rain, or Neither is true; your evidence supports .4 credence in Snow, .3 credence in Rain, and .3 credence in Neither. In this case, your rational credence in Snow is higher than either your rational credence in each of Rain and Neither. So, your evidence here favors snow over its alternatives.

Second, like Guess Calibration, Evidential Calibration is about *educated guesses*, cast in all-or-nothing terms. As we understand it, to make a rational educated guess that P, it is not necessary to also form an all-out *belief* that P. Rather, making an educated guess that P is compatible with suspending judgment about P, or having some intermediate degree of credence in P. Depending on your theory of belief, it may be compatible with believing \sim P as well. What is necessary to rationally guess that P is that your credence in P is higher than your credence in any of the alternatives to P. To get the idea, consider what would be rational to guess in a forced-choice situation. Let's return to the weather examples above. If you were forced to choose between the options carved up as {Rain, \sim Rain}, and your credence in Rain was below .5, you would make an educated guess that it will not rain. If you were forced to choose between the options carved up as {Rain, Snow, Neither} and your credence in Rain was higher than your credence in either Snow or Neither, you would choose Rain. (It is this kind of case – one in which your options are more finely partitioned – in which it might be rational to guess that P even if you are more confident of \sim P than of P.)¹⁴ Your educated guess should thus correspond to the option that you have the highest credence in.

Third, what is the *expected reliability* of one's educated guess that P? If your guess corresponds to the option you have the highest credence in, the *reliability* of your guess is the probability that you will assign the highest credence to the option that is *correct*, or true – it is the probability, that you would get the answer right, if you had to choose. Since rational agents are often uncertain about their reliability, we are concerned here with rational *expected*

¹⁴ So it can sometimes be rational for you to make a particular guess, even if you think that guess is probably wrong. This is just a normal consequence of uncertainty. A student taking a multiple choice test might pick the option that is only marginally more plausible to her than each of the others. Although she is rational in picking the option that has the best shot at being right, and can even defend her choice on that basis, she might still rationally expect to get the question wrong. This situation does not seem Moore-paradoxical to us. On the SATs, for example, it doesn't seem at all problematic to say something like: "Well, out of A, B, C, and D, C seems the most plausible. I'm about 40% sure that's the answer. So probably C is wrong. But if it is wrong, I have no idea which answer is right: each of A, B, and D are less plausible than C." Importantly, this kind of situation will arise in cases where your choices range over several options – not just P and \sim P.

reliability: the weighted average of different possibilities for how reliable you could be, as assessed by your own lights. We will follow White’s modal understanding of “reliability” and “expected reliability” – like Guess Calibration, Evidential Calibration is concerned with your expected propensity to guess correctly. Your expected propensity to guess correctly that P in a given situation depends on at least two features of that situation: it depends both on your cognitive capacities – your ability to evaluate your first-order evidence – *and* on the type of first-order evidence that’s available to you.¹⁵ (Note that your expected reliability may well come apart from your *actual* reliability – your actual propensity to guess correctly. This will be often be the case when you have misleading evidence. In fact, this is exactly what happens in Anton’s situation. Anton’s expected reliability is .4, while his actual reliability (given that the mushrooms he took were just regular dried porcini) is much higher.) Our expected reliability is generally constrained by the type of evidence that is available to us and by how good we are at evaluating it. Evidential Calibration says that your credence in the option that your evidence favors should match the expectation of how reliable your educated guess will be.¹⁶

That’s all we’ll say now about the content of Evidential Calibration. Why should we think it’s right? First, Evidential Calibration satisfies Desideratum 1: it says that one’s rational beliefs should be determined, in part, by the bearing of one’s first-order evidence. So, in Anton’s case, Anton should be more confident in B than in any of the alternatives because his first-order evidence favours B over A, C, and D. This means that Evidential Calibration

¹⁵ In Anton’s case, propensity to guess right “in a given situation” means something like this: looking at the particular type of test results, while under the influence of Sam’s magic mushrooms. The relevant situation here must be understood in way that is “independent of” or “prior to” his particular evidence or reasoning about which dose to give his patient. (If we included Anton’s first-order evidence and allowed it to have its usual epistemic role, it would seem that Anton should remain confident in his guess on the basis of that evidence.) The intuitive thought here is clear enough, although spelling out the appropriate “independence principle” is a delicate job. We will assume here that the job can be done. See Christensen [2007], [2009], [2011], Elga [2007], and Vavova [ms], among others, for further discussion of these issues. If no good independence principle can be formulated, this will raise challenges for any account (including Guess Calibration) on which higher-order evidence rationally affects first-order beliefs.

¹⁶ Could your rational expected reliability about some topic ever exceed the reliability of your evidence? (For example: if your first-order evidence only indicated whether P with .6 reliability, could the expected reliability of your making an educated guess that P, on the basis of E, be .8?) It’s not obvious how such situations could arise. Suppose that your higher order evidence indicates that you are perfectly reliable at evaluating your evidence. That means that whenever your evidence supports P to .6, you will correctly evaluate it as supporting P to .6. But then, you should rationally expect that your reliability about P will *match* the reliability of your first-order evidence. In fact, it’s hard to see how we would fill out a story so that you could rationally expect to be *more* reliable than your first-order evidence. What would explain this greater expected reliability? A guardian angle that guides you to the right answer on the occasions that your evidence turns out to be misleading? Or a knack for extra sensory perception? And would these be scenarios on which your expected reliability exceeds your first-order evidential support – or would they be scenarios in which you simply have some *additional* source of first-order evidence (the guardian angle, or ESP)?

gives the right verdicts in the cases that created problems for Guess Calibration. First, consider the case in which your evidence rationalizes .8 credence in P and you have no reason to doubt your reliability. Nevertheless, you slip up and become .8 confident in $\sim P$ instead. Unlike Guess Calibration, Evidential Calibration does *not* yield the verdict that you are fully rational: Evidential Calibration says that if your evidence rationalizes .8 credence in P, the rational response is to be .8 confident in P. Evidential Calibration also captures the asymmetry between Anton and Anna (Anton's colleague who mistakenly concluded that A, rather than B, is the right dose). Unlike Guess Calibration, Evidential Calibration says that Anna is not fully rational. This is because both Anton and Anna should be .4 confident in the dose that their shared first-order evidence *actually supports*; if Anna ends up with .4 confidence in A, she is not fully rational.

Second, Evidential Calibration satisfies Desideratum 2, by allowing higher-order evidence to rationally influence our first-order attitudes. In doing so, Evidential Calibration is in line with verdicts various "conciliatory" or "equal weight" positions have given regarding how we should respond to particular kinds of higher-order evidence, like peer disagreement.¹⁷

Finally, Evidential Calibration explains a general and plausible intuition about the relationship between one's rational attitudes about one's own reliability and one's rational first-order attitudes: it explains why someone who is rationally certain that she has responded to her evidence rationally is generally entitled to be more opinionated than someone who has doubts about her own rationality. This is because according to Evidential Calibration, your rational credence in P should match the expected reliability of your making an educated guess that P. And, as we have seen, the expected reliability of making an educated guess that P will typically depend on both the strength of your first-order evidence *and* your own ability to assessing the evidence rationally. Thus, suppose again that you are looking at some meteorological data, D, and you know that given evidence like D, it tends to rain 75% of the time. Furthermore, you are completely confident that you are identifying and accommodating your meteorological evidence correctly. How confident should you be that it will rain?

¹⁷ See, especially, Elga [2007] on how one should accommodate evidence of disagreement:

Equal weight view: Upon finding out that an advisor disagrees, your probability that you are right should equal your prior conditional probability that you would be right. Prior to what? Prior to your thinking through the disputed issue, and finding out what the advisor thinks of it. Conditional on what?

On whatever you have learned about the circumstances of the disagreement.

Evidential Calibration accommodates Elga's principle, but it goes farther: it is formulated more generally, and it gives an account of what to believe on one's *total* evidence, making the contribution of one's first-order evidence clear.

Intuitively, you should be .75 confident. And Evidential Calibration gives us exactly this verdict. That's because since you are certain that you will evaluate your evidence correctly, you should expect that you will *make an educated guess* that P if and only if the evidence *in fact favors* P. If your educated guess perfectly tracks your evidence, your expected reliability of guessing that P should be the same as the degree of your evidential support for P. Your credence that it will rain, given that you make an educated guess that it will rain based on your evidence, should be the same as your credence that it will rain given your first-order evidence alone – that is, .75. In these situations, Evidential Calibration reduces to the principle that your rational credence in P should match the probability of P given the first-order evidence.

What if you have reason to think that you *don't* always evaluate the evidence correctly – that perhaps sometimes you guess that P even when your evidence favors some alternative to P? In general, evaluating our evidence rationally tends to lead us to the truth; usually our evidence isn't systematically misleading. If you guess that P, but your evidence favors \sim P, it is likely that your guess is false. Reasons to doubt our rationality often give us reasons to suspect that we will be unreliable: so, evidence that we're *irrational* will usually lower our expected reliability as well.¹⁸ And, according to Evidential Calibration, when we have reason to doubt our reliability, we should revise confidence in our first-order beliefs. Evidential Calibration therefore makes sense of the intuition that when we suspect that we might be irrational, we should leave ourselves a doxastic margin for error.¹⁹

¹⁸ The exact relationship between your expected reliability and these other factors may not be straightforward or simple, and we cannot address the question in full here. But we can say some general things about this relationship. First, if you are rationally confident that you are perfectly reliable at rationally assessing your evidence, your expected reliability should just equal the strength of that evidence. Could you rationally regard yourself as *anti-reliable* with respect to some question – i.e. could your expected reliability be significantly lower than chance? According to Egan and Elga (2005), the answer is “no”. On pain of incoherence, an agent with decent access to her own beliefs must assign low probability to the claim that she is an anti-expert about the subject matter. Thus, there are, arguably, independent constraints on the lower bound of your expected reliability.

¹⁹ Before moving on, we'd like to address a potential worry. Calculation is a case in which Anton has fairly targeted higher-order evidence: it suggests that he is unreliable in making judgments about a particular subject matter, namely which dose is the right one for his patient. But what about cases in which our higher-order evidence is more general – suggesting that Anton is bad at *all kinds* of reasoning, including making educated guesses about his reliability, making educated guesses about the reliability of those guesses, and so on? One might worry that in a case like this, our view will require Anton to apply Evidential Calibration over and over again at every level, and there will be no stable answer as to what his rational credence in dose B should be.

Whether such general defeaters are possible is controversial. (See Roush [2009] for skepticism about this.) Supposing that such situations can come about, one possible response is to say that, in such cases, the rational requirements are undefined. (See Lasonen-Aarnio [forthcoming] for discussion of this kind of solution.) This possibility is also open to us in other paradoxical cases (such as self-ascribed anti-expertise, mentioned in fn. 18). Though we find this option attractive, we won't defend it at length here. But notice that the problem of general defeat – if it is possible – is not a special worry for Evidential Calibration. Rather, it's a worry for any

4. Meeting Desideratum 3

We saw that Level-Splitting, which straightforwardly accommodates both first-order and higher-order evidence, faces serious problems in its rejection of Desideratum 3: it allows one's rational first-order and higher-order attitudes to be in conflict. If Evidential Calibration also accommodates both kinds of evidence, won't it have to reject Desideratum 3 as well?

The aim of this section is to look more closely at what it takes to satisfy Desideratum 3. What does it mean for first-order and higher-order attitudes to conflict, and why is this conflict irrational? First, we will take a step back and look at one of the original motivations for Desideratum 3: the thought that it's irrational to have high confidence in both "P" and "my evidence doesn't support P". We argue that Evidential Calibration can vindicate this thought, and helps us rule out a certain class of bad combinations of first-order and higher-order attitudes that Level-Splitting permits; this gives us reason to think that Evidential Calibration does respect Desideratum 3, after all. In the last part of this section we bring up a remaining puzzle for Evidential Calibration and suggest a response.

4.1 "P, but my evidence doesn't support P"

In cases like Calculation, Level-Splitting views say that it can be rational to have high confidence in both "P" and "my evidence doesn't support P". This combination of attitudes strikes us as clearly irrational, and a clear violation of Desideratum 3. But it would be nice if we could say more about what's wrong with it, and why it's not *merely* counterintuitive on its face. In this section we provide such an argument.²⁰

To see what's wrong with believing (or having high confidence in) these Moore-paradoxical combinations, let's look back at Level-Splitting. In a single case like Calculation, Level-Splitting says that it can be rational to have high confidence in both "P" and "my evidence doesn't support P". This is implausible on its face, but further problems emerge when we consider what the view will say if cases like Calculation happen again and again.

view that takes higher-order evidence seriously. Dealing with this type of case will therefore pose a challenge for any view that respects Desideratum 2.

²⁰ Though many authors assume that this combination of attitudes is irrational, direct arguments for this claim are surprisingly hard to come by. See Christensen [2007] for an argument that one should not be certain of some logical truth, while being less than certain of one's own rationality in deducing the logical truth; see also White [2009] and [ms]; and [reference omitted]. The arguments in this section are expanded upon in [reference omitted]. See Elga [forthcoming] for an independent, but complementary, proposal regarding when level-tension is and is not rational.

Suppose that after talking to Sam, Anton calculates the doses for a number of other patients. His reasoning yields the doses P, Q, R, S, etc. Since we've assumed that Anton is not in fact rationally impaired, let's suppose that these calculations are all correct, and that in each case Anton has responded rationally to his evidence. But Anton doesn't know this; he is confident that he is under the influence of Sam's mushrooms.

Level-Splitting will say that in each case, Anton should become highly confident of the dose that his first-order evidence supports. At the end of the long sequence of cases, then, Anton should be rationally highly confident of each of the following:

P: P is the correct dose for patient 1.

Q: Q is the correct dose for patient 2.

... etc.

If Anton is keeping track, he can survey his prescription history and conclude the following:

P1: I got all of these calculations (P, Q, R, S, ...) right. My track record is amazing!

But according to Level-Splitting, Anton should also be rationally highly confident in:

P2: Only 40% of P, Q, R, S, ... are supported by my evidence. I'm only 40% reliable at reaching rational educated guesses.

P1 and P2, together, seem to license some odd-looking reasoning. Anton can wonder: "How strange! Since I expected to be 40% reliable, I thought I would only get the right answer about 40% of the time. But I'm doing much better than that – I've had a much better success rate than 40%. I must be more reliable than I thought I was. What could explain my excellent performance?" Initially, Anton might just think that he's getting lucky. But as he continues to compile a better and better track record, the divergence between his actual track record and the track record he rationally expected, given his expected reliability, will become more and more remarkable. The chances that he would irrationally guess correctly so many times in a row are very slim.

So, reflecting on his surprisingly good track record, Anton should be able to conclude that there is some other explanation for why his track record systematically exceeds what he would have expected. Specifically, he should conclude:

C: The mushrooms must not have affected me! I'm actually highly reliable.

This is not a rational way of coming to believe that one's educated guesses are reliable.

Crucially, it's problematic that Anton can go through this reasoning without ever performing

an independent check on whether his doses were right, or on whether he had in fact been drugged. This is bootstrapping.²¹

Anton's doxastic state in Calculation – having high confidence in P, and having high confidence that his evidence doesn't support P – licenses bad reasoning. So there must be something wrong with Anton's doxastic state, and the particular type of tension between Anton's first-order and higher-order attitudes.

4.2 No problem for Evidential Calibration

We now have an argument against a certain type of level-tension, permitted by Level-Splitting: it licenses bootstrapping. In this section, we'll see that this argument gives Evidential Calibration a substantial advantage.

First, we should note that unlike Level-Splitting, Evidential Calibration does *not* allow bootstrapping. According to Level-Splitting, Anton could bootstrap because he could acquire a track record that was much better than he would have predicted, given his expected reliability. But Evidential Calibration straightforwardly rules out this possibility.²² According to Evidential Calibration, one's confidence in any proposition P *cannot* rationally exceed one's expected reliability about P. Suppose Anton follows Evidential Calibration while calculating several patients' doses in a row. If his expected reliability is only 40% in each case, Anton's confidence in each of P, Q, R, and S should be only .4. So his apparent track record will be no better than he would have predicted.

We can see why Level-Splitting faces a problem here, and why Evidential Calibration does not, by comparing Level-Splitting to simple reliabilism about justification. On a simple reliabilist view, one's belief that P is justified if and only if it was produced by a reliable process: one need not know, or even believe, that the process is reliable. But if this is right, someone wondering whether one of her (in fact reliable) belief-forming processes is reliable could very easily find out that it is by bootstrapping. All she would need to do is form some beliefs using that process, note (without an independent check) that all of these beliefs are true and were formed using the process, and conclude via induction that the process must be reliable.²³

²¹See White [2009] and [ms], as well as Christensen [2007], for discussion of similar points. See Vogel [2000] and Cohen [2002], e.g., for more general discussion of bootstrapping.

²²See also White [2009] for discussion of how Guess Calibration prevents rational bootstrapping.

²³See, e.g., Vogel [2000] and Cohen [2002].

Simple reliabilism sanctions bootstrapping because it lacks a principle like Evidential Calibration. Even when you have reason to think that a certain (in fact reliable) belief-forming process is *unreliable*, simple reliabilism says that you can rationally use the process anyway. More sophisticated reliabilist views try to rule out rational bootstrapping by building in a defeat condition: Goldman, for instance, holds that evidence that a process is *unreliable* can defeat our justification in beliefs formed through that process.²⁴ On these more sophisticated views, when you have positive reason to believe that a process is unreliable, you cannot rationally maintain the beliefs produced via that process. So, according to these views, you cannot rationally acquire a good track record when you rationally expect that your track record will be bad.

But by rejecting a principle like Evidential Calibration, Level-Splitting explicitly rules out this kind of defeat – at least when the belief-forming process you are doubting is your own rationality. Because of this, Level-Splitting also allows bootstrapping.

Evidential Calibration can help us understand what’s wrong with the kind of level-tension that Level-Splitting recommends in cases like Calculation. And it can also help us understand why this kind of tension seemed so odd to begin with. Since evidence that we are irrational is often also evidence that we are unreliable, we should in be less confident in P when we have reason to suspect that our educated guesses about P are unsupported by our evidence.

4.3 Desideratum 3 as Evidential Calibration?

We set out, in this section, to see whether Evidential Calibration satisfies Desideratum 3; we worried that, in taking account of both first-order and higher-order evidence, Evidential Calibration would end up in just as bad a position as Level-Splitting. By now we have seen enough to draw some preliminary conclusions. First, we saw an argument against certain paradigmatic violations of Desideratum 3. While Level-Splitting runs into problems with this argument, Evidential Calibration does not. We also saw that Evidential Calibration can help us understand why Level-Splitting had these problems, and gave an explanation for the intuitive oddness that violations of Desideratum 3 produce.

All of this suggests that Evidential Calibration does, in fact, meet Desideratum 3. Evidential Calibration gives us a plausible picture of the relationship between epistemic

²⁴See Goldman [1986]. See Smithies [2012] for further (skeptical) discussion of Goldman’s no-defeaters clause.

levels that respects and vindicates our core intuitive judgments. In fact, we might take Evidential Calibration itself to embody the right way of making Desideratum 3 more precise.

4.4 Rational level-tension for Evidential Calibration?

Given what we have seen so far, there is good reason to think that Evidential Calibration rules out problematic level-conflicts, and therefore satisfies Desideratum 3. But Evidential Calibration does not rule out *all* level-conflicts as rational.

We originally observed that it seems irrational to have high confidence in “P, but my evidence doesn’t support P”. Someone in this epistemic state has high credence in P, but takes her credence to be much higher than the credence warranted by her evidence. Generalizing this thought, it is tempting to think that we weaker types of mismatch are irrational, too. For example, it seems odd for someone to think that her credence in P is *too high*, even if she is unsure about which lower credence it would be rational to adopt. (There may be a number of different candidate credences that she considers as possibilities, but that are on balance lower than the credence in P that she currently has.²⁵)

We can capture this tempting line of thought with the following, more precise principle:

Rational Reflection: An agent’s first-order credence in P should match her expected rational credence in P.²⁶

Your *expected rational credence* in a proposition is average of all of the candidate credences that you think might be rational, weighted by how likely you take it to be that each one is rational. When you think that your credence in P is too high, but are unsure about which lower credence it would be rational to adopt, your credence in P comes apart from your expected rational credence in P. By prohibiting this kind of mismatch, Rational Reflection places some limits on which combinations of first- and higher order states can be rational: for example, it will prevent you from having rational high confidence in something like “P, but my evidence supports \sim P”, or “P, but my evidence probably supports a lower level of confidence than the one I have.”

²⁵ Some authors would argue that in cases like this, we should have imprecise or “mushy” credences. We won’t discuss this possibility here, as it will not affect our arguments.

²⁶ See Christensen [2010] gives a formal definition of this principle:

Rational Reflection: $Cr(A|Pr(A) = n) = n$

where Cr is your credence in A, and Pr is the function describing the ideally rational credence for you to have in A. Strictly speaking, our version of Rational Reflection is entailed by Christensen’s formal account, but does not entail it; however, this distinction won’t make a difference for our purposes. See Christensen [2010] and Elga [forthcoming] for further discussion of this principle.

Rational Reflection is an initially plausible account of tension between epistemic levels. But it is incompatible with Evidential Calibration.²⁷ To see why, let's return to Calculation.²⁸ Recall that Evidential Calibration said that Anton is rationally required to be .4 confident that dose B was right, and that furthermore, this is the *only* rational response to Anton's situation. Let's suppose that Anton becomes .4 confident of dose B, as Evidential Calibration suggests. Let's also suppose that Anton, being a self-conscious follower of Evidential Calibration himself, believes there is only one rational response to his total evidence: to have .4 confidence in whatever dose his first-order evidence actually favors. Putting those pieces together, Anton cannot be rationally fully confident that his .4 confidence in B is rational. If the mushrooms make his educated guesses less reliable, then they must do so by impairing his ability to assess his evidence. Anton should be less than certain that the dose that he guessed to be right (dose B) was in fact the dose favored by his evidence.

But if Anton's evidence doesn't favor dose B, it must favor some other dose. (This kind of evidence, we can assume, always supports one dose over the others.) If Anton's evidence favors some other dose, then his credence in dose B is irrationally high. So Anton should think that his credence in dose B might be too high. He also knows that in his situation, the highest credence that could be rational for him to have in *any* dose is .4. So according to Evidential Calibration, Anton might rationally think: "My credence in dose B might be rational, or it might be too high. But it's definitely not too low."²⁹

Anton should be uncertain whether his first-order doxastic attitude is rational, and furthermore, he should think that on balance, *if* he's irrational, there is a particular direction in which he is erring. According to Evidential Calibration, then, Anton's credence in dose B should *not* match his expected rational credence in dose B.

²⁷ For the record, Rational Reflection is also incompatible with Level-Splitting. Since Level-Splitting licenses high confidence in "P, but most likely my evidence doesn't support P", e.g., it allows one's rational credence in P to be much higher than his expected rational credence in P.

²⁸ Our discussion here follows Christensen [2010] and Elga [forthcoming]. Our Calculation case is meant to parallel the case of the unmarked clock discussed in those papers, and in Williamson [notes a] and [notes b], with at least one salient difference: in the clock case, we are asked to imagine someone who is uncertain about what her evidence *is*, but knows what is *supported* in various evidential situations. In Calculation, Anton knows what his evidence *is* and is unsure what it supports.

²⁹ Suppose, for instance, that in addition to having .4 credence in the right dose, Anton is rationally required to divide the remainder of his credence evenly over the wrong answers. So, if B is the right dose, Anton should have .2 credence in each of A, C, and D; if A is the right dose, Anton should have .4 credence in A, and .2 credence in B, C, and D, etc. If Anton knows that this is what Evidential Calibration recommends, he can know his rational credence in dose B will never be higher than .4 – it is .4 if B is right, and .2 otherwise.

Rational Reflection gives us a prima facie plausible way to spell out what it is for epistemic levels to conflict. But as a rational requirement, it turns out to be incompatible with Evidential Calibration. So we face a choice: give up Evidential Calibration, or accept that some instances of tension between epistemic levels can be rational.

We think that the second option is an attractive one. First of all, we have already seen that, in accepting some rational level-tension, Evidential Calibration does not put us in the same dire situation as Level-Splitting. Rather than rejecting any systematic rational relation between epistemic levels, as Level-Splitting does, Evidential Calibration says that there is such a relationship, and helps us understand what that relationship is. Second, in accepting Rational Reflection as an epistemic norm, we risk running into the problems we have already seen for views that give up Desideratum 1 or 2: it may be hard to adopt Rational Reflection without facing (one version of) those problems as well. Third, Rational Reflection rules out nearly all cases of rational uncertainty about what one should believe.³⁰ This may seem like the wrong result – especially given situations like Calculation, in which it is plausibly rational to be unsure about what one’s evidence supports.

And finally: while Rational Reflection is one way of predicting and explaining our intuitive judgments, it would be nice to have an additional argument showing whether these intuitive judgments are right, and whether Rational Reflection is the right way to explain them. We have such an argument in favor of Evidential Calibration. Without a corresponding argument to bolster intuitions in favor of Rational Reflection, we should tentatively conclude that, while mismatch between one’s credence in P and one’s expected rational credence in P is odd, it is not always irrational. Adopting Evidential Calibration yields a surprising conclusion about which combinations of first-order and higher-order attitudes we can rationally hold. But this conclusion is not, we think, too surprising to accept.³¹

5. Conclusion

We started by motivating three desiderata for a plausible account of higher-order evidence: we said that rational beliefs should reflect the bearing of (1) first-order evidence *and* (2) higher-order evidence, but that (3) one’s first-order and higher-order attitudes should not be in tension with one another. It seems impossible to meet all three desiderata in a

³⁰ One kind of uncertainty it permits is cases where, for example, one is unsure whether one’s credence is too high or too low, but takes either possibility to be equally likely. See Christensen [2010] for further discussion.

³¹ Elga [forthcoming] draws a similar conclusion regarding Rational Reflection, but approaches the issue from a different angle.

straightforward way. But it is possible to strike a balance, and Evidential Calibration does: it holds us responsible for accommodating our first-order evidence correctly, while allowing us to be rationally uncertain about what to believe. Evidential Calibration also provides us with a principled account of how epistemic levels should interact: why our first-order and higher-order attitudes should generally be aligned, and why it is sometimes problematic when they diverge.

Bibliography

- Christensen, David. [2007 a] "Epistemology of Disagreement: the Good News". *The Philosophical Review*, volume 116, no. 2, pp. 187-217.
- [2007 b] "Does Murphy's Law Apply in Epistemology? Self-Doubt and Rational Ideals", *Oxford Studies in Epistemology* 2, 3-31.
 - [2010] "Higher-Order Evidence". *Philosophy and Phenomenological Research* 81: 1 pp. 85 - 215.
 - [2010] "Rational Reflection". *Philosophical Perspectives*, volume 24: Epistemology (Blackwell, Oxford) pp. 121-140.
 - [2011] "Disagreement, Question-Begging and Epistemic Self-Criticism", *Philosopher's Imprint* 1, no. 6.
- Coates, Allen. [2012] "Rational Epistemic Akrasia". *American Philosophical Quarterly* volume 1, no. 2, 113-124.
- Cohen, Stewart. [2002] "Basic Knowledge and the Problem of Easy Knowledge". *Philosophy and Phenomenological Research* 65 (2):309-39.
- Egan, Andy and Elga, Adam. [2005] "I can't believe I'm stupid". *Philosophical Perspectives* 19 (1): 77-93.
- Elga, Adam. "The Puzzle of the Unmarked Clock and the New Rational Reflection Principle". Forthcoming in *Philosophical Studies*.
- [2007] "Reflection and Disagreement". *Noûs*, volume XLI, no. 3, pp. 478-502.
 - [ms] "Lucky to be Rational".
- Feldman, Richard. [2005] "Respecting the Evidence". *Philosophical Perspectives*, volume 19: Epistemology (Blackwell, Oxford) pp. 95-119.
- Goldman, Alvin [1986]. *Epistemology and Cognition*, Harvard University Press.
- Hazlett, Allan. [2012] "Higher-Order Epistemic Attitudes and Intellectual Humility". *Episteme* 9 (3): 205-223.
- Kelly, Thomas. [2010] "Peer Disagreement and Higher Order Evidence". In Richard Feldman and Ted Warfield (eds.), *Disagreement* (OUP): 11-174.
- [2005] "The Epistemic Significance of Disagreement". In John Hawthorne & Tamar Gendler (eds.), *Oxford Studies in Epistemology*, Volume 1. Oxford University Press.

- Lasonen-Aarnio, Maria. "Higher-Order Evidence and the Limits of Defeat". Forthcoming in *Philosophy and Phenomenological Research*.
- Roush, Sherrilyn. [2009] "Second-Guessing: A Self-Help Manual". *Episteme* 6 (3): 251-268.
- Smithies, Declan. [2012] "Moore's Paradox and the Accessibility of Justification". *Philosophy and Phenomenological Research* 85 (2): 273-300.
- Titelbaum, Michael. [ms] "In Defense of Right Reason".
- Vavova, Ekaterina. [ms] "Irrelevant Influences".
- Vogel, Jonathan. [2000] "Reliabilism Leveled". *Journal of Philosophy* 97 (11):602-623.
- Weatherson, Brian. [ms] "Do Judgments Screen Evidence?"
- Wedgwood, Ralph. [2011] "Justified Inference". *Synthese* 189 (2): 1-23.
- White, Roger [2009]. "On Treating Oneself and Others as Thermometers". *Episteme* 6 (3): 233-250.
- [ms] "Disrespecting the Evidence".
- Williamson, Timothy [2000]. *Knowledge and its Limits*. Oxford: Oxford University Press.
- [notes a] "Improbable Knowing".
- [notes b] "Very Improbable Knowing".