

# Can your total evidence mislead about itself?

Alex Worsnip

Forthcoming in Mattias Skipper Rasmussen & Asbjørn Steglich-Petersen (eds.), *Higher-Order Evidence: New Essays*, Oxford University Press

*Penultimate draft; please cite published version where possible*

## 1. Introduction: what would it be for total evidence to mislead about itself?

It's fairly uncontroversial that you can get misleading evidence about your evidence. Suppose the clues at the crime scene do not support believing  $p$ , but then some expert tells you (in a rare lapse of judgment) that the clues do support believing  $p$ . Then you have misleading "higher-order evidence" about your "first-order evidence" (and what it supports believing). More generally, we can define misleading higher-order evidence as follows:

**Misleading Higher-Order Evidence.** A case of misleading higher-order evidence is one where *either*:

- (a) (i) Your first-order evidence supports some doxastic attitude  $D$  toward  $p$ , but (ii) you have higher-order evidence that your first-order evidence does not support  $D$  toward  $p$ ; *or*
- (b) (i) Your first-order evidence does not support some doxastic attitude  $D$  toward  $p$ , but (ii) you have higher-order evidence that your first-order evidence supports  $D$  toward  $p$ <sup>1</sup>

What is more controversial is whether such a phenomenon can ever result in a situation whereby one has misleading *total* evidence about one's total evidence. That is:

**Misleading Total Evidence (about Total Evidence).** A case of misleading total evidence (about total evidence) is one where either:

- a) (i) Your total evidence supports some doxastic attitude  $D$  toward  $p$ , but (ii) your total evidence supports believing that your total evidence does not support  $D$  toward  $p$
- b) (i) Your total evidence does not support some doxastic attitude  $D$  toward  $p$ , but (ii) your total evidence supports believing that your total evidence supports  $D$  toward  $p$

Such cases, are there are such, are ones where one's total evidence misleads about what *it itself* supports.

These are general characterizations of misleading higher-order evidence and misleading total evidence, but I am going to focus on the particular kind of misleading higher-order evidence that I

---

This paper is greatly indebted to a conversation with Miriam Schoenfield, and an objection that she made to my (2018). Though it was initially introduced for the purposes of this objection, the basic idea behind the model in this paper is due to her. I am also grateful to Nilanjan Das, David James Barnett, Ram Neta, and Jim Pryor for helpful discussion, and to an anonymous referee and the editors of this volume for extremely helpful comments.

<sup>1</sup> You may notice that, under both (a) and (b), clause (i) refers to evidential support for *attitudes*, whereas clause (ii) refers to evidence for *propositions*. (Moreover, in the definition of misleading total evidence below, both clauses (i) and (ii) refer to evidential support for attitudes.) This is intentional; I refer you to §2 below to understand my terminology.

mentioned in the opening paragraph: that is, cases where your first-order evidence does not support believing  $p$ , but your higher-order evidence supports believing that your first-order evidence does not support believing  $p$ . The question is: can such a case amount to a case of misleading *total* evidence, where your total evidence does not support believing  $p$ , but your total evidence supports believing that your total evidence supports believing  $p$ ? If we can show that it does, then we have shown that at least one kind of case of misleading total evidence, defined in the more general above way, is possible.

## 2. Straightening out some terminology

In trying to make progress on this question, we need to be clear about our terminology. One source of confusion in the literature is that not all participants use the term ‘support’ in the same way. First, some talk of the evidence supporting *propositions*, while others talk of the evidence supporting *attitudes* toward those propositions. Secondly, to say that the evidence “supports” something might mean either that it supports it to *some* degree, or that it supports it to some relevantly *sufficient* degree. I think there is probably more than one permissible way of talking here, but there are some considerations that affect choice of terminology. First, we want to allow ourselves the capacity to say that in some cases, the attitude toward a proposition  $p$  that is best supported by one’s total body of evidence is that of *suspending judgment* about whether  $p$ . Secondly, though, within such cases, we want to be able to distinguish between those in which one’s total body of evidence, on balance, supports  $p$  to *some* extent (just not enough to support believing it), and those in which it doesn’t. We don’t want to mistakenly suggest that whenever the evidence on balance favors  $p$  over not- $p$ , it warrants believing  $p$ .

In light of that, I will talk (and have already been talking, in §1) in the following way. Evidential support for *propositions* fundamentally comes in degrees. We can talk about *how strongly* a body of evidence supports a proposition  $p$ . When I say that a body of evidence supports  $p$ , without a qualifier as to how strongly or weakly it supports it, I mean only that it supports  $p$ , on balance, *at least somewhat better than it supports not- $p$* . This is still compatible with its supporting  $p$  (over not- $p$ ) rather weakly – not strongly enough to warrant believing  $p$ , rather than suspending judgment.

In addition to this, however, I will *also* talk of evidential support for attitudes. When I talk of a body of evidence supporting an attitude, I mean that given that body of evidence (taken on its own), the attitude is warranted. This is an all-or-nothing matter, and is not something that comes in degrees. But plausibly, it is systematically related to degreed evidential support for propositions. Plausibly, for a body of evidence to support believing  $p$  just for it to support  $p$  to some (perhaps situationally variable) sufficient degree (to warrant belief). It’s also natural to say that what it is for a body of evidence to support suspending judgment about whether  $p$  is for it neither to support  $p$  to a sufficient degree (to warrant belief), nor to support not- $p$  to a sufficient degree. But this latter claim will have to be complicated if there are permissive cases where the evidence permits either believing  $p$  or

suspending judgment about whether  $p$ .<sup>2</sup> For the sake of simplicity, I will ignore permissive cases in what follows, though the core of my argument does not turn on assuming that there are no such cases.

In a sense, then, ‘support’ means something different when I talk about support for *propositions* than it does when I talk about support for *attitudes*. Perhaps it would be good to have two different words instead of using ‘support’ in these two different ways. But my meaning can always be discerned by whether I am talking of support for propositions or for attitudes.

I also need to clarify what *does* the supporting. On the way I’m talking, a “body of evidence” can be one’s *total* evidence, or it can be a subset of one’s total evidence (at the limit, it could be a single item, e.g., a single proposition). So, we can ask both whether one’s total evidence supports  $p$ , and whether some particular subset of one’s evidence supports  $p$ . Slightly less obviously, I think that we can also talk either about whether one’s total evidence supports an attitude toward  $p$ , or about whether a subset of one’s total evidence supports an attitude toward  $p$ .<sup>3</sup> Indeed, notice that, as I defined “misleading higher-order evidence” in §1 above, it refers explicitly in clauses (a.i) and (b.i) to the support provided for *attitudes* by subsets of one’s evidence. It might seem a little unclear what this talk comes to, given my stipulations about what it means to talk of evidential support for attitudes. But I propose to understand it as follows: to say that a subset of one’s evidence supports a doxastic attitude toward  $p$  is to say that, given that subset of one’s evidence, and *bracketing all one’s other evidence*, the attitude is warranted. (This is what the “taken on its own” clause in my definition of support for attitudes was supposed to indicate.) Put counterfactually, if that helps: *were* that subset of one’s evidence to exhaust one’s total evidence, that doxastic attitude toward  $p$  would be warranted.<sup>4</sup> Note that this does not require that the subset of one’s evidence makes that doxastic attitude warranted even given the presence of one’s other evidence.

However, I do want a word to refer to the stronger feature, of a subset of one’s evidence, of *making* it the case that one’s total evidence supports a doxastic attitude. I will say that a subset of one’s evidence is “dispositive” with respect to a particular doxastic attitude when, given the other evidence one has, its presence makes it the case that one’s total evidence supports this attitude. Given my terminology, then, a subset of your evidence can support a doxastic attitude without being dispositive with respect to that attitude: this is so when, were it to exhaust your total evidence, it would warrant belief, but the other evidence you have is such that overall, your total evidence does not (sufficiently) support believing this proposition.

I will characterize higher-order evidence, somewhat vaguely, as evidence *about* one’s evidence, or what one’s evidence supports, or how one has responded to one’s evidence. I won’t try to

---

<sup>2</sup> If this is so, we might try saying that for the evidence to support suspending judgment about whether  $p$  is for it not to support  $p$  so strongly that *only* believing  $p$  is warranted, and likewise for not- $p$ . Or we might want to introduce notions of ‘decisive’ and ‘sufficient’ support, where decisive support for an attitude entails that *only* that attitude is warranted.

<sup>3</sup> In light of these points, I’ll avoid using the language of ‘*pro tanto* support’, because it’s not clear whether it refers to support *to a degree that may not be sufficient to warrant belief*, or whether it refers to support *by a part, rather than the whole, of one’s evidence*. These two statuses are orthogonal and should not be confused.

<sup>4</sup> This counterfactual is tricky to manage when we are talking about *higher-order* evidence, though. Since higher-order evidence is (at least often) evidence *about* one’s first-order evidence, it’s hard to envisage a situation where a particular body of higher-order evidence exhausts one’s evidence. Still, I think there’s still a good sense in which we can evaluate whether the support provided by the higher-order evidence for a proposition would be enough to warrant belief if we at least *bracket* the force of the first-order evidence.

characterize exactly what this aboutness relation comes to. I assume that we have a good enough grip on it to classify some evidence as first-order and some as higher-order: for example, the clues at the crime scene are first-order evidence, whereas testimony about what the clues support is higher-order evidence. Even if it turned out that no principled or completely sharp distinction between first-order evidence and higher-order evidence could be drawn, though, this would not be a problem for my argument. In principle, I could expunge all reference to “first-order” and “higher-order” evidence and just talk about the clues and the testimony, arguing that ultimately this package of evidence can lead to a situation where one’s *total* evidence, consisting of both the clues and the testimony, misleads about itself. With that point made, I will continue to talk of “first-order” and “higher-order” evidence in what follows.

One point about the meaning of ‘higher-order evidence’ bears stressing, however. It is yet another frequent source of confusion in the literature that the adjectives “higher-order” and “first-order” can attach either to different bodies of *evidence*, or to different *propositions* (and attitudes that have those propositions as their objects). A higher-order proposition would be a proposition like *the evidence supports believing p*, where the corresponding first-order proposition is *p*. It might be tempting to characterize the higher-order evidence as the evidence that evidentially bears on the higher-order proposition, and the first-order evidence as the evidence that evidentially bears on the first-order proposition. But this would be a mistake. As we’ll see shortly, most philosophers (on both sides of the debate about misleading total evidence) admit that higher-order evidence, such as testimony about whether the evidence supports believing *p*, *can* be at least *some* evidence for *p* itself. (And, slightly more arguably, first-order evidence, such as the clues, may provide some evidence concerning higher-order propositions about what they themselves support.) So the higher-order evidence cannot be distinguished from the first-order evidence just by what propositions it bears on.<sup>5</sup> Maybe there is a sense in which higher-order evidence bears more *directly* on the relevant higher-order proposition than it does on the relevant first-order proposition, but it can still bear on the latter in a perfectly good sense – or at least, this should not be ruled out by terminological fiat.

### 3. (Re)framing the debate

With this terminology in hand, we can reframe the debate. In particular, we can state two different ways in which one might have misleading higher-order evidence without having misleading total evidence (about total evidence). First, one’s higher-order evidence might be dispositive *both* with respect to believing the (higher-order) proposition that one’s total evidence supports believing *p*, *and* with respect to believing the (first-order) proposition *p* itself. Many philosophers have thought that at least usually, “evidence of evidence is evidence” – that is, slightly more precisely, evidence that one has evidence for *p* is itself evidence for *p*. If this is so, then even when one lacks first-order evidence that supports believing *p*, the misleading higher-order evidence that one has might itself be dispositive

---

<sup>5</sup> Given this, the ‘aboutness’ relation referred to in the previous paragraph had better not be that of evidential bearing. I assume it is something stronger. One stronger way of characterizing it, assuming that evidence is propositional, would be to say that the higher-order evidence is evidence *about* one’s evidence in that its propositional content ineliminably refers to one’s evidence. But I leave it open whether this is exactly right.

with respect to believing *p*. In such a case, one's total evidence would support believing that one's total evidence supports believing *p*, and one's total evidence *would* support believing *p*. Thus, though we are in a case of misleading higher-order evidence, we would not be in a situation of misleading total evidence (about total evidence). Call such cases **doubly-dispositive cases** – since they are cases where the misleading higher-order evidence is dispositive *both* with respect to the higher-order belief *and* the first-order belief.

Secondly, one's higher-order evidence might be dispositive *neither* with respect to believing the proposition that one's total evidence supports believing *p*, *nor* with respect to believing the proposition *p* itself. In such a situation, one's higher-order evidence would provide *some* support for the proposition that one's total evidence supports believing *p* – perhaps even enough such that, taken on its own, it supports believing this proposition – but, given the presence of the first-order evidence, not enough to be dispositive with respect to believing this proposition. Remember that in the case of misleading higher-order evidence that I'm focusing on, one's first-order evidence does not support believing *p*. Maybe this fact carries with it some kind of rational self-evidence that is not defeated by higher-order evidence such as testimony. In such a case, one's total evidence would not support believing *p*, nor would one's total evidence support believing that one's total evidence supports believing *p*. Thus, again, though we are in a case of misleading higher-order evidence, we are not in a position of misleading total evidence (about total evidence). Call such cases **non-dispositive cases** – since they are cases where the misleading higher-order evidence is dispositive *neither* with respect to the higher-order belief *nor* the first-order belief.

Faced with cases of misleading higher-order evidence, some philosophers try to block the possibility of misleading *total* evidence (about total evidence) by claiming that all cases of misleading higher-order evidence are doubly-dispositive cases;<sup>6</sup> others by claiming that all cases of misleading higher-order evidence are non-dispositive cases.<sup>7</sup> But, in my view, it is more promising for the denier of misleading total evidence (about total evidence) to make the weaker, and more plausible, claim that all cases of misleading higher-order evidence are *either* doubly-dispositive cases *or* non-dispositive cases. That is: the misleading higher-order evidence *may* be dispositive with respect to the higher-order belief, or may not be; the crucial claim is simply that *if* it is dispositive with respect to the higher-order belief, it is also dispositive with respect to the first-order belief.

The challenge for those who think that there *can* be misleading total evidence (about total evidence), then, is to show that this is mistaken. In other words, they must show that there can be **singly-dispositive cases**: cases where the misleading higher-order evidence is dispositive with respect to the higher-order belief, but not the first-order belief.

Let me spell out exactly why singly-dispositive cases of misleading higher-order evidence, if there are such cases, are instances of misleading total evidence about total evidence. Begin with a case in which your first-order evidence does not support believing *p*, but you have misleading higher-order evidence that your first-order evidence does support believing *p*. As I said above, the possibility of such a case, so far, should be uncontroversial. Now, if the higher-order evidence is singly-dispositive, then it is dispositive with respect to the belief that your total evidence supports *p* – meaning that this

---

<sup>6</sup> Cf., e.g., Feldman (2005), Bergmann (2005), and (more qualifiedly) Horowitz (2014).

<sup>7</sup> Cf., e.g., Titelbaum (2015).

belief is now supported by your total evidence. At the same time, if the higher-order evidence is singly-dispositive, then it is *not* dispositive with respect to believing *p*. So, since your first-order evidence doesn't support believing *p*, and your higher-order evidence isn't dispositive with respect to believing *p*, your total evidence doesn't support believing *p*. Thus, any singly-dispositive case of misleading higher-order evidence is one where your total evidence supports believing that your total evidence supports believing *p*, but where your total evidence doesn't support believing *p*. That is, it is a case of misleading total evidence about total evidence.

I've tried before (though not using the same terminology) to argue for the possibility of singly-dispositive cases (Worsnip 2018). But though I think they are on the right track, I'm not fully happy with the arguments I've provided previously. In this paper, I am going to explore a slightly more formal argument, using a simple mathematical model. I believe this argument poses a significant challenge to those who deny the possibility of misleading total evidence, though as I will acknowledge, it still falls short of an incorrigible proof.

As a final preliminary to my core argument, let me reiterate that those who affirm the possibility of misleading total evidence (about total evidence) *do* tend to acknowledge that misleading higher-order evidence has *some* evidential bearing on the relevant first-order belief.<sup>8</sup> I concur here – evidence of evidence *is* (typically, *some*<sup>9</sup>) evidence – and my argument will not attempt to show otherwise. Indeed, it may even be that misleading higher-order evidence is *often* dispositive with respect to the relevant first-order belief. My view is simply that misleading higher-order evidence can *sometimes* be dispositive with respect to a higher-order belief, while not being dispositive with respect to the corresponding first-order belief (that is, that it is sometimes singly-dispositive). If that is so, then misleading total evidence (about total evidence) is possible.

The plan for the remainder of this paper is as follows. In §4, I will argue for the following claim: misleading higher-order evidence is often *stronger* evidence for the relevant higher-order proposition than it is for the relevant first-order proposition. Call this the Comparative Claim. I think that the Comparative Claim can be established fairly decisively, using the model that I provide. In §5, I will explore how we might get from this claim to the claim that there are singly-dispositive cases: cases where the higher-order is dispositive with respect to the relevant higher-order belief without being dispositive with respect to the relevant first-order belief. This step might seem trivial, but it turns out not to be. Nevertheless, I will suggest that it is still fairly hard to resist this step of the argument. And if one accepts that step of the argument – if one accepts that there are singly-dispositive cases – it follows, as I've already shown in this section, that misleading total evidence about total evidence is possible.

#### 4. Arguing for the Comparative Claim

---

<sup>8</sup> Cf., e.g., Kelly (2010); Pryor (2013: 99-100); Worsnip (2018); Lasonen-Aarnio (forthcoming).

<sup>9</sup> See Fitelson (2012) for a case where evidence that there is evidence for *p* does not itself provide any evidence for *p*. But Fitelson explicitly says (translating to my terminology) that this only works for evidence that is not dispositive. By his own lights, his case is not one where one has dispositive evidence that one's evidence supports believing *p*. In fact, he says that he suspects that when one does have such dispositive evidence that one's evidence supports believing *p*, *this fact is* always evidence for *p* (*ibid.*: fn. 2).

As I stated it above, the Comparative Claim was vague. I will begin by arguing for the following claim;

Comparative Claim<sub>1</sub>: There are cases where one's higher-order evidence supports the higher-order proposition that one's *first-order* evidence supports believing  $p$  more strongly than it supports the first-order proposition  $p$ .

I will then argue from that claim to another, subtly different Comparative Claim, which is the one that is needed to subsequently argue to the possibility of misleading total evidence about total evidence.

a) *Arguing for Comparative Claim<sub>1</sub>*

Let us work with a more concrete case. I've used this example before, in Worsnip (2018).<sup>10</sup> In this case, Miss Marple, an expert detective, and her niece Mabel, a novice, visit a murder scene, where they survey the (first-order) evidence – the clues. (Suppose they gather exactly the same first-order evidence as each other, and they both know this.) This (first-order) evidence does not support any positive belief about who committed the crime. However, Miss Marple – who is generally an expert about what the evidence supports – makes an uncharacteristic mistake – and declares that the (first-order) evidence supports believing that the vicar did it. In fact, the first-order evidence supports suspending judgment about whether the vicar did it; in fact, let us suppose,<sup>11</sup> it doesn't even on balance support the proposition that the vicar did it.

Clearly, this case is one where Mabel has misleading higher-order evidence about what the first-order evidence supports. However, as the case was just stated, there seem to be few principled grounds on which to arbitrate whether the case is either a doubly-dispositive case, a non-dispositive case, or a singly-dispositive case – and it is only if it is a singly-dispositive case that it is a case of misleading total evidence (about total evidence). But remember that to begin with, I only want to argue that, at least given a few eminently possible stipulations, the Miss Marple case establishes the Comparative Claim – that Miss Marple's testimony supports the relevant higher-order proposition *more strongly* than it supports the relevant first-order proposition.<sup>12</sup>

Let  $F$  be the (first-order) proposition that the vicar did it,  $H$  be the (higher-order) proposition that the first-order evidence supports believing that the vicar did it, and  $T$  be the proposition that Miss Marple testifies that  $H$ . So, what we are investigating is whether there is a version of the Miss Marple case such that  $\Pr(H|T) > \Pr(F|T)$ .

The value on the left-hand side – the probability that the evidence supports believing the vicar did it, conditional on Miss Marple's testifying that the evidence supports believing the vicar did it – is, essentially a matter of the reliability of Miss Marple's testimony (assuming that Mabel knows how reliable Miss Marple is). So, for example, if Mabel knows that Miss Marple is 90% reliable with respect

---

<sup>10</sup> In turn, it's based on other similar cases in Coates (2012) and Horowitz (2014).

<sup>11</sup> And this is a further supposition; see §2 for a reminder of the terminology.

<sup>12</sup> Thanks to the editors of this volume for helpful suggestions about how to present the formal details of the foregoing argument.

to relevantly similar matters – then  $\Pr(H|T) = 0.9$ .<sup>13</sup> This is so because Miss Marple’s testimony is directly about whether H.

The value on the right-hand side is a little trickier to obtain, since Miss Marple doesn’t directly testify about whether F, but only about whether H. But it can be done. Given that Miss Marple testifies that the evidence supports believing that the vicar did it – that she testifies that H – there are exactly two possibilities under which the vicar *did* do it.

- 1) Miss Marple is *right* that the evidence supports believing that the vicar did it, and he did in fact do it.
- 2) Miss Marple is *wrong* that the evidence supports believing that the vicar did it, but as a matter of fact the vicar nevertheless did do it.

Let’s assume for now, to simplify, that if it is already known whether H, T (i.e., Miss Marple’s testifying that H) does not change the probability of F – in other words, T and F are probabilistically independent conditional on whether H. (I’ll return to this assumption in a moment.) Given this assumption, the probability of the first possibility, given Miss Marple’s testimony, is in effect the probability that the evidence does support believing that the vicar did it, given that Miss Marple testifies that it does, multiplied by the probability that the vicar did it, given that the evidence supports believing that he did. More formally, it is:  $\Pr(H|T).\Pr(F|H)$ .

The probability of the second possibility, given Miss Marple’s testimony, is in effect the probability that the evidence does *not* support believing that the vicar did it, given that Miss Marple testified that it does, multiplied by the probability that the vicar did it, given that the evidence does *not* support believing that he did. More formally, it is:  $\Pr(\sim H|T).\Pr(F|\sim H)$ .

By adding these the probabilities of each of the two possibilities together, then, we can get the probability that the vicar did it, conditional on Miss Marple’s testifying that the evidence supports believing that the vicar did it. In other words:

$$\Pr(F|T) = \Pr(H|T).\Pr(F|H) + \Pr(\sim H|T).\Pr(F|\sim H)$$

Given this, we can now find out whether  $\Pr(H|T) > \Pr(F|T)$ , for any case, as long as we have the values of the following three parameters:

- A)  $\Pr(H|T)$ <sup>14</sup>
- B)  $\Pr(F|H)$
- C)  $\Pr(F|\sim H)$

---

<sup>13</sup> Of course, there may be other evidence bearing on H, as I’ll discuss in §4 below; here we are just considering the probability of H conditional on T alone.

<sup>14</sup> Since  $\Pr(\sim H|T) = 1 - (\Pr(H|T))$ , we don’t need to separately know the value of the former, even though it occurs in the above formula for  $\Pr(F|T)$ .



And indeed, there are many values of (A), (B) and (C) such that  $\Pr(H|T) > \Pr(F|T)$ .<sup>15</sup> For example, suppose  $\Pr(H|T) = 0.9$ ,  $\Pr(F|H) = 0.9$ , and  $\Pr(F|\sim H) = 0.5$ . Less formally: given that Miss Marple says the evidence supports believing the vicar did it, there's a 0.9 probability that the evidence does support believing he did it; given that the evidence supports believing he did it, there's a 0.9 probability that he did do it; and given that the evidence doesn't support believing he did it, there's a 0.5 probability that he did do it. There's no reason to think that these values of (A), (B) and (C) aren't co-possible. And given these values:

$$\begin{aligned}\Pr(H|T) &= 0.9 \\ \Pr(F|T) &= (0.9).(0.9) + (0.1).(0.5) = 0.86\end{aligned}$$

Thus, in this version of the case,  $\Pr(H|T) > \Pr(F|T)$ .

Note also that we can set possible (and, indeed, plausible) values for A, B, and C that create a much bigger gulf between  $\Pr(H|T)$  and  $\Pr(F|T)$ . In general, as parameter A [i.e.  $\Pr(H|T)$ ] goes up, parameter B [i.e.  $\Pr(F|H)$ ] goes down, and parameter C [i.e.  $\Pr(F|\sim H)$ ] goes down, the gulf will get bigger. For example, suppose that  $\Pr(H|T) = 0.95$ ,  $\Pr(F|H) = 0.8$ , and  $\Pr(F|\sim H) = 0.2$ .<sup>16</sup> Then:

$$\begin{aligned}\Pr(H|T) &= 0.95 \\ \Pr(F|T) &= (0.95).(0.8) + (0.05).(0.2) = 0.77\end{aligned}$$

It follows that Comparative Claim<sub>1</sub> is true: we have a case where one's higher-order evidence supports the higher-order belief that one's first-order evidence supports believing p more strongly than it supports first-order belief in p.

Let me now return to the simplifying assumption that I made, namely that T and F are probabilistically independent conditional on whether H. In general, this seems a plausible assumption. If it's already known for certain that H (i.e. that the first-order evidence supports believing T), then it seems that T – Miss Marple's saying that H – doesn't add further evidential support for T. However, it might be challenged. Consider a case where it's known that  $\sim H$  – that the first-order evidence does not support believing that the vicar is guilty.  $\sim H$  does not, in and of itself, say that the first-order evidence does not, on balance, support the proposition that the vicar is guilty to *some* degree; it merely says that the first-order evidence does not support this proposition to a *sufficient* degree to warrant believing that the vicar is guilty. Given this, it might be contended that Miss Marple's testimony that H is still some evidence that F, even if we already know that  $\sim H$ . After all, perhaps Miss Marple is more likely to have made a slight error – declaring that the first-order evidence sufficiently supports

---

<sup>15</sup> Not every possible set of values for X, Y and Z will yield the desired result, but again, we only need *one* case where they do yield this result to secure the *possibility* of misleading total evidence (about total evidence).

<sup>16</sup> How could the probability that the vicar did it, conditional on the evidence not supporting believing that he did it, be anything other than 0.5? Easily. First, the logical space may be partitioned such that the unconditional probability that the vicar did it is low (e.g. there may be lots of different suspects). Secondly, conditionalizing on the evidence *not* supporting believing that the vicar did it rules out the worlds in which the evidence *supports* believing that the vicar did it, leaving *both* the worlds in which the evidence support suspending judgment on whether the vicar did it, *and* the worlds in which the evidence supports believing that the vicar *did not* do it. So in general, one would expect the probability that the vicar did it, conditional on the evidence *not* supporting believing that he did it, to be lower than its unconditional probability.

believing that the vicar is guilty when it does, on balance, provide some evidence for that proposition, but not enough – than to have made a more serious error – declaring that the evidence sufficient supports believing that the vicar is guilty when it doesn’t even support that proposition on balance to *any* degree. If that is so, even knowing that Miss Marple’s testimony is false, it still provides *some* support for the proposition that the vicar is guilty – that is, for F.<sup>17</sup> That is, T and F are not independent conditional on  $\sim H$ .

This is a good point, but we could make further stipulations about the case so as to neutralize it. Suppose Miss Marple is always in one of two states. In her sober state, she’s *infallible* with respect to what the evidence supports. But when she’s in the (unusual) state of having had three or more glasses of sherry, she’s *no better than chance* with respect to what the evidence supports. (Unfortunately, let’s add, Miss Marple gives no outward signs when she’s in the latter state.) Given those stipulations, the case in which Miss Marple is wrong that the clues support believing that the vicar did it are guaranteed to be one in which she’s drunk, and thus no better than chance. So, in this version of the case, it is no longer true that, conditional on  $\sim H$ , T is evidence for F. And so the simplifying assumption is true at least in this case, which is enough to show that Comparative Claim<sub>1</sub> is true.

Moreover, even when we relax the simplifying assumption, it’s still plausible that we’ll get some cases where  $\Pr(H|T) > \Pr(F|T)$ . In cases where T raises the probability of F even conditional on  $\sim H$ , we need in our calculations to replace parameter (C) –  $\Pr(F|\sim H)$  – with what we could call parameter (C\*) –  $\Pr(F|\sim H \& T)$ . Recall that in my initial demonstration, I gave two examples where  $\Pr(H|T) > \Pr(F|T)$ . In the first, I supposed that  $C = 0.5$ ; in the second, I supposed that  $C = 0.2$ ; even the former supposition was enough to get the desired result. (See fn. 15 on why the value of C can fall well below 0.5.) In examples like the second, where the value of C is very low, the value of C\* will plausibly still be quite low, lower (for example) than the value of C in the first example. And that, again, will be enough to deliver the result that  $\Pr(H|T) > \Pr(F|T)$ .

(b) *From Comparative Claim<sub>1</sub> to Comparative Claim<sub>2</sub>*

So far, we have Comparative Claim<sub>1</sub>, which, to repeat, is:

Comparative Claim<sub>1</sub>: There are cases where one’s higher-order evidence supports the higher-order proposition that one’s *first-order* evidence supports believing p more strongly than it supports the first-order proposition p.

Contrast this with:

Comparative Claim<sub>2</sub>: There are cases where one’s higher-order evidence supports the higher-order proposition that one’s *total* evidence supports believing p more strongly than it supports the first-order proposition p.

---

<sup>17</sup> Thanks to an anonymous referee for making this point.

The argument directly establishes Comparative Claim<sub>1</sub>, rather than Comparative Claim<sub>2</sub>, because as we defined it, H is the (higher-order) proposition that the *first*-order evidence supports believing that the vicar did it. Indeed, this is not accidental, for this is what Miss Marple testifies about in our case: she reports on what the clues support, not (directly) about what the totality of the evidence, including her own testimony, supports.

However, for the argument for singly-dispositive cases (and thus, for the possibility of misleading total evidence about total evidence), we want Comparative Claim<sub>2</sub>. After all, the endgame is to show that one's total evidence can support believing falsehoods *about itself*, which requires us to focus on the higher-order belief that one's *total* evidence supports believing p.

Fortunately, for any normal case, it is a short step from Comparative Claim<sub>1</sub> to Comparative Claim<sub>2</sub>. Recall that H is the proposition that the first-order evidence supports believing that the vicar did it. Let H\* be the proposition that the total evidence supports believing that the vicar did it. Now, fortunately, in any normal case, Miss Marple's testimony will be at least as strong evidence for H\* as it is for H. Why is that so? Consider what Mabel's total evidence consists in: the first-order evidence, *plus* Miss Marple's testimony itself. Now, very plausibly, Miss Marple's testimony (that the first-order evidence supports believing that the vicar did it) is *some* evidence that the vicar did it – and Mabel knows this. And at the very least, Miss Marple's testimony is certainly not evidence that the vicar *didn't* do it. So the difference between the first-order and total evidence here is just that the latter also includes a piece of evidence that very plausibly speaks in favor of believing, and manifestly doesn't speak *against* believing, that the vicar did it.

Given this, it seems that to whatever extent Miss Marple's testimony supports H – the proposition that the *first*-order evidence supports believing that the vicar did it – it will also support H\* – the proposition that the *total* evidence supports believing that the vicar did it – to at least the same extent. In order for it not to do so, Miss Marple's testimony – that the first-order evidence supports believing that the vicar did it – would have to indicate *of itself* that it speaks against believing that the vicar did it (while all the while still being evidence that the *first*-order evidence supports believing that the vicar did it). It's hard to see how this could possibly be so. And even if it were possible,<sup>18</sup> it is evidently not the norm.

So, it seems safe to assume that in at least the vast majority of cases, Miss Marple's testimony is at least as strong evidence for H\* as it is for H. But if Miss Marple's testimony supports H more strongly than it supports F (as established in arguing for Comparative Claim<sub>1</sub>), and it supports H\* at least as strongly as it supports H (as I've just argued), then it supports H\* more strongly than it supports F. That is, Comparative Claim<sub>2</sub> follows after all.

Thus, I conclude that Comparative Claim<sub>2</sub> is true. There can be cases where misleading higher-order evidence supports a higher-order belief about one's total evidence more strongly than it supports the corresponding first-order belief.

## 5. From the Comparative Claim to the possibility of misleading total evidence

---

<sup>18</sup> Note, by the way, how odd it would be for someone who is trying to *limit* the capacity of evidence to mislead about itself to affirm this possibility.

But how do we get from Comparative Claim<sub>2</sub> to the possibility of misleading total evidence? The general shape of the strategy is as follows. Comparative Claim<sub>2</sub> asserts that there are cases where misleading higher-order evidence more strongly supports a higher-order belief of the form *my total evidence supports p* than it supports the first-order belief *p*. Now, not *all* of these cases will be singly-dispositive cases. It might be that even though the misleading higher-order evidence supports the higher-order proposition more strongly than the first-order proposition, it is still strong enough to be dispositive with respect to believing both propositions, or weak enough to be dispositive with respect to believing neither. However, the thought is this: there should be cases at the margins, where given that the misleading higher-order evidence supports the higher-order proposition more strongly than the first-order proposition, it is strong enough to be dispositive with respect to the higher-order belief, but not with respect to the first-order belief. Such cases are singly-dispositive, and as I already demonstrated in §3, any singly-dispositive case is a case of misleading total evidence (about total evidence).

As stated, the above argument is suggestive but somewhat vague. One tempting way to precisify it might be as follows. Assume, as I've already suggested, that there is a threshold of degressed evidential support for a proposition above which belief in that proposition is evidentially supported in an on-off sense.<sup>19</sup> For any reasonable threshold, we will be able to find values of parameters (A), (B) and (C) such that  $\Pr(H|T)$  – and therefore, by the argument of §4(b),  $\Pr(H^*|T)$  – exceeds the threshold, but  $\Pr(F|T)$  doesn't. For example, if the threshold is somewhere between 0.86 and 0.9, then the original set of values for the parameters are such that  $\Pr(H^*|T)$  exceeds the threshold but  $\Pr(F|T)$  doesn't. If it is anywhere between 0.77 and 0.95, then the second set of values for the parameters are such that  $\Pr(H^*|T)$  exceeds the threshold but  $\Pr(F|T)$  doesn't.

However, it is too quick to conclude from this that the case is one where belief in  $H^*$  is supported by the *total* evidence but belief in  $F$  isn't. For this is a function of the evidential support provided for (or against)  $H^*$  and  $F$  not only by Miss Marple's testimony (T), but also by the remainder of the evidence: the first-order clues (call their totality C). Again, it is tempting to argue as follows. By stipulation of the case, C does not support  $F$ . Therefore, C can't possibly make up for the shortfall in support that T provides to  $F$  compared to that it provides for  $H^*$ . Thus, T must in principle be capable of being dispositive with respect to believing  $H^*$  without being dispositive with respect to believing  $F$ .

Again, though, the argument is too quick. It overlooks the following inconvenient point: while C merely *fails to provide evidence for*  $F$ , arguably C *actively provides evidence against*  $H^*$ . Remember what  $F$  and  $H^*$  say.  $F$  is the proposition that the vicar did it. By stipulation of the case, C (the totality of the first-order clues), on balance, supports suspending judgment about  $F$ .  $H^*$ , though, is the proposition that the total evidence supports believing  $F$ . Suppose that the clues enjoy at least some degree of rational self-evidence: that is, they support, at least to some degree, whatever proposition is true about the evidential support that *they themselves* provide (or do not provide). In that case, C, to at least some

---

<sup>19</sup> Again, we need not assume the threshold is the same for all propositions and situations, though the argument as stated does assume that for any one situation, the threshold for  $H^*$  and that for  $F$  is the same. This might be questioned. But even if the threshold for the higher-order belief is higher than that for the first-order belief, I've already shown that the gulf in how strongly Miss Marple's testimony supports the two propositions can be quite large, and so it still seems likely that  $\Pr(H^*|T)$  could exceed the threshold while  $\Pr(H|T)$  doesn't.

degree, supports the proposition that the first-order evidence does not support believing F – which is just the negation of H\*. That is, C provides at least *some* evidence against H\*.

Given this, the opponent of misleading total evidence (about total evidence) might say that, to whatever extent that T supports H\* more strongly than it supports F, this is cancelled out by the fact that C actively provides evidence against H\*, while merely failing to provide evidence for F. Thus, it can never be the case that one's *total* evidence supports H\* more strongly than it supports F. So, even at the margins, it will never be the case that one's total evidence on balance supports believing H\*, without on balance supporting believing F.

This is an important objection, and I do not have a knock-down response to it. Let me begin with a rather defensive point to clarify the dialectic here. It is still very much a conjecture on the objector's part that the evidence that C provides against H\* is *always* enough to make up for the difference in evidential support that T provides for H\* as opposed to F. It's unclear what justifies this conjecture. Remember that the defender of the possibility of misleading total evidence only needs there to be *some* cases in which the total evidence supports believing H\* but does not support believing F. Given that the gulf between the support that T provides to H\* and the support that it provides to F may in some cases be quite big, and given that (as I'll also shortly argue) the evidence that C provides against H\* may be quite weak, arguably the natural default position is still to hold that there is nothing to rule out the possibility of misleading total evidence about total evidence. Because of the vagaries of exactly how strongly C tells against H\*, the argument here is going to fall short of an incorrigible proof. Still, I think it is a real step forward, in making a case for misleading total evidence, to have shown that T is stronger evidence for H\* than it is for F. If T was always *just as strong* evidence for F as it is for H\*, then – especially in light of the fact that C provides some evidence against H\* – it really would be hard to see how T could ever result in a situation of misleading total evidence about total evidence. As things are, though, the opponent of misleading total evidence is now on the back foot: to defend her view, she is forced into *a priori* conjectures, that must hold as strict universal generalizations, about the relative strengths of different evidential relations, and it's not clear what the argument for such conjectures would be.

So, now, why is it that the evidence that C provides against H\* may be relatively weak? I will discuss three reasons. The first is that, as clarified in the last section, H\* is a proposition about what the *total* evidence supports believing. So, even if C provides even fairly weighty evidence about its *own* epistemic weightiness (or, in this case, lack of), it does not follow that it is particularly weighty evidence about what the *total* evidence – including Miss Marple's testimony – supports. Of course, by indicating that it itself doesn't support the proposition that the vicar did it, it provides *some* evidence about what the total evidence supports, since it is part of the total evidence. But this need not be weighty. And this is so especially in light of the fact that it only indicates of itself that it doesn't support H\*, not that it is strong evidence against it. For virtually any proposition p, presumably there is some part of your evidence that isn't evidence for p; learning this will often be very insignificant as evidence that your *total* evidence doesn't support p.

A second feature of the case that makes C's evidential force against H\* weak *for Mabel* is the fact that Mabel is a novice detective.<sup>20</sup> Though Mabel is directly acquainted with the clues, and though

---

<sup>20</sup> See Kappel (forthcoming: esp. §3.2) for some points related to this second reply.

the clues have a particular property, namely that of supporting suspending judgment about whether the vicar is guilty, this doesn't entail that Mabel is directly acquainted with *the fact that the clues (have the property of) support(ing) suspending judgment about whether the vicar is guilty*. (Compare: you can be directly acquainted with a piece of music, and that music can have certain properties – e.g. the duration of the piece, the exact intervals between notes, etc – without your being directly acquainted with the fact that the piece of music has those properties.) In order for the clues themselves to provide Mabel with evidence about their own properties (e.g. that they support suspending judgment about whether the vicar is guilty), she has to have some capacity to register the properties of the clues. (Again, just as, in order for the music to provide one with evidence about its own properties, one has to have some capacity to register to properties of the music.)

Now, we can grant that Mabel does have this capacity to *some* degree. But it is also true that the *strength* of the evidence that the clues provide with respect to their own properties is sensitive to the *degree* of Mabel's capacity to register these properties; to how good she is at doing so. (Once more, the *better* you are at registering the properties of a piece of music, the weightier the evidence that the music provides with respect to its own properties.) And since Mabel is a novice detective, her capacity to register the properties of the clues is not especially refined. That weakens the evidence that the clues supply to Mabel with respect to their own properties. Because the clues only provide Mabel with evidence about their own properties *via* Mabel's capacity to recognize these properties, shortfalls in this capacity attenuate the strength of the evidence they provide.<sup>21</sup> This is why, in coming to a view about H\*, some other expert (say, Sherlock Holmes) will be justified in giving more weight to the clues than a novice (like Mabel) will be – as compared to the weight given to Miss Marple's (countervailing) testimony.

This point can be sharpened by attending to the distinction between the *balance* and the *weight* of evidence (see Joyce 2005). One's evidence can be *balanced* quite strongly against p, in the sense that what evidence one does have all points against p, while still not being very *weighty* evidence against p, in the sense that it wouldn't take much countervailing evidence to swing one's total balance of evidence in the opposite direction. If we think in terms of credences, then the *balance* of the evidence determines what one's credence in p should be, while the *weight* of the evidence determines how *resilient* that credence should be in response to future evidence. Importantly, there can be differences in the weight of two different evidential situations that nevertheless are balanced in the same way. (Compare a situation where there is no evidence concerning p whatsoever with one where there is a large amount of evidence for p and an equally large amount against it. In both situations, the evidence is precisely evenly balanced, such that one's credence should be 0.5. But in the latter, the evidence has much

---

<sup>21</sup> Note that this isn't to make the (somewhat) more controversial claim that the evidence provides support for first-order propositions only to the extent that one has the capacity to recognize that it provides support for those first-order propositions. The (generalized version of the) claim I'm making here is just that the evidence provides support for higher-order beliefs *about which first-order attitudes it itself provides support for* only to the extent that one has the capacity to recognize that it provides support for those first-order attitudes. The latter claim ought to be uncontroversial, even among those who dispute the former claim. Even if my being acquainted with the clues, plus the clues supporting a first-order proposition, suffices for my possessing the clues as evidence for that first-order proposition (regardless of my capacity to recognize them as such), this doesn't change the fact that if I can't recognize them as evidence for the first-order proposition, I won't possess evidence for the higher-order proposition *that the clues support the first-order proposition* (or belief in it).

greater weight than in the former – such that one’s credence should be much more *resilient*, that is, it ought to take a lot more future evidence to significantly shift one’s credence away from 0.5.)

Now, compare the evidential situation of Mabel and Sherlock. For both of them, C is some evidence against H\*. Moreover, C may in *one* sense be equally strong evidence against H\* for the two of them, if by that we mean that the *balance* of the evidence, when C is the *only* evidence, is the same for the two of them. But if what I’ve suggested is right, C is less *weighty* for Mabel than it is for Sherlock. That is, even if C should (on its own) make Mabel’s credence in H\* low, this low credence should nevertheless not be very resilient, relative to other potential evidence, such as Miss Marple’s testimony. By contrast, Sherlock’s credence should be at least somewhat more resilient in the face of Miss Marple’s testimony (though not completely so). If we focus only on the balance of the evidence provided by C, and not its weight, we miss this difference between Mabel and Sherlock, and so we miss a crucial sense in which the evidential force of C against H\* *for Mabel* is weak.

Thirdly, and finally, to the extent that C does provide evidence against H\*, it is possible that at least some of this evidential support is *defeated* by Miss Marple’s testimony (T). Following Pollock (1986), it’s standard to distinguish “rebutting” and “undercutting” defeaters. On one way of drawing that distinction, a “rebutting” defeater is one that defeats one’s justification for believing p by offering positive evidence that not-p, while an “undercutting” defeater is one that defeats one’s justification for believing p by in some way casting doubt on one’s existing evidence for p. Importantly, undercutting defeat can take place even in cases where one’s existing evidence for p is *in fact* (or at least, *would be*, in the absence of the defeater) good evidence for p. That is, even *misleading* evidence that some piece of evidence E is not good evidence for p can undermine E’s (would-be) support of p: the phenomenon is not restricted to case where E only *apparently* supports p, and one’s higher-order evidence correctly indicates that it doesn’t.

Given these characterizations, it is plausible to claim that T is *both* a rebutting *and* an undercutting defeater for the evidence that C provides against H\*. It is rebutting because T is itself evidence *for* H\*; but it is also an undercutting defeater because it directly casts doubt on the claim that C is itself evidence *against* H\*.<sup>22</sup> So in *addition* to providing evidence for H\*, T may at least partially disarm or disable the evidential force that C has against H\*. If that is so, C’s force against H\* will be even weaker than it would have been without the presence of T.

It might be thought odd that I am appealing to defeat to help me in this way. In general, defeat-talk is used more often by those who *resist* the possibility of misleading total evidence.<sup>23</sup> The idea is that when higher-order evidence makes it rational to believe that some first-order attitude is irrational,

---

<sup>22</sup> Not all cases of rebutting defeat are also cases of undercutting defeat. Take an ordinary case where you have two fairly weighty pieces of evidence, for and against a proposition. For example, suppose that you’re trying to figure out whether Jones will be at the party. On one hand, there will be dancing at the party, which is evidence that Jones will be there (since Jones loves dancing); on the other hand, there will be singing at the party, which is evidence that Jones won’t be there (since Jones hates singing). Though these two pieces of evidence cut against each other in the sense that they support opposite conclusions, and certainly either piece of evidence makes it rational to be less confident than it would be rational to be if one had only the other piece of evidence, still neither piece of evidence casts any doubt on the other piece of evidence’s status as evidence for the opposite conclusion. (At least, this is so if we don’t develop the case such that all of Jones’ enjoyment of dancing goes away in the presence of singing, or that all of Jones’ hatred of singing goes away in the presence of dancing.)

<sup>23</sup> See e.g. Skipper Rasmussen (this volume); González de Prado Salas (ms.). On the other side, Lasonen-Aarnio (2014) accepts the possibility of misleading total evidence but resists (to some degree) the phenomenon of (higher-order) defeat.

it also defeats the justification (or evidential support) for the first-order attitude itself, thus avoiding the possibility of misleading total evidence. However, as I've suggested before (Worsnip 2018: 31-32), this strategy is very hard to extend to cases – like the one that we are presently discussing – where the first-order attitude supported by the first-order evidence is one of suspending judgment. If the clues (taken on their own) supported *believing* that the vicar did it, or at least on balance supported the proposition that the vicar did it, perhaps misleading higher-order evidence (that the clues do not support believing that the vicar did it) could defeat the evidence that the clues would have otherwise provided that the vicar did it. But in the case where the clues (taken on their own) do not tell one way or the other as to whether the vicar did it, it's unclear what there is for misleading higher-order evidence (that the clues *do* support believing that the vicar did it) to “defeat”, at least in the undercutting sense. It's not like this misleading higher-order evidence can make the *clues themselves* actually support believing that the vicar did it.

Thus, the (rather complex) dialectical situation is this. I imagined my opponent pointing out that while the clues (C) merely fail to support F, they actively provide evidence *against* H\* (such that even if T supports H\* more strongly than it supports F, the total evidence might still fail to support H\* more strongly than it supports F). My (third) reply was that to the extent that C provides evidence against H\*, this evidential support may be at least partially defeated by T. At the same time, C's *failure to support F* isn't the sort of thing that can be defeated (in the undercutting sense) by T. So the defeating effect that T has will tend to neutralize the asymmetry of support that C has with respect to F and H\*: it will bring us back toward a situation where C does not tell for or against either F *or* H\*.

Drawing together the three replies I've given here, I conclude that C is fairly weak evidence against H\*, and that whatever evidence it does provide is arguably undercut, at least in part, by T. This makes it hard to believe that C's evidential force against H\* will always make up for the shortfall in support that T provides for F compared with its greater support for H\*. And so, I think we still have reason to think that, in principle, there can be situations where one's total evidence supports H\* without supporting F – that is, in situations of misleading total evidence about total evidence.

## 6. Conclusion

In closing, let me briefly say a bit about why this matters.<sup>24</sup> Most epistemologists assume that, at least in some sense, you ought to follow your evidence: that is, to believe what your total evidence supports, and refrain from believing those things that your total evidence does not support. If there can be cases of misleading total evidence about total evidence, then this entails that sometimes, it is both the case that you ought to believe that your evidence supports believing p, but yet also that you ought not to believe p. But if you have this combination of states, you are “epistemically akratic”. To put it another way, your doxastic states exhibit a kind of incoherence across levels. It is also widely thought that such akratic or incoherent combinations of attitudes are irrational; or, even more strongly, that it is not *possible* to sustain such states in full reflective awareness of them. So, if there can be misleading total evidence about total evidence, we must revise some piece of orthodoxy: either the claim that you always ought to follow your evidence, or that epistemic akrasia is always irrational, or that there cannot

---

<sup>24</sup> Here I'm very briefly summarizing the puzzle laid out in Worsnip (2018) and in Lasonen-Aarnio (forthcoming).



be dilemmatic or “tragic” situations where the normative injunction to follow your evidence and the normative injunction not to be akratic come into conflict. Thus, the possibility of misleading total evidence bears crucially on utterly foundational issues in the theory of epistemic rationality.

## References

- Bergmann, M. (2005). ‘Defeaters and Higher-Level Requirements,’ *The Philosophical Quarterly*, 55/220: 419-436.
- Coates, A. (2012). ‘Rational Epistemic Akrasia,’ *American Philosophical Quarterly*, 49/2: 113-124.
- Fitelson, B. (2012). ‘Evidence of evidence is not (necessarily) evidence,’ *Analysis*, 72/1: 85-88.
- González de Prado Salas, J. (ms.) “Dispossessing Defeat.”
- Horowitz, S. (2014). ‘Epistemic Akrasia,’ *Noûs*, 48/4: 718-744.
- Joyce, J.M. (2005). ‘How Probabilities Reflect Evidence,’ *Philosophical Perspectives*, 19: 153-178.
- Kappel, K. (forthcoming). ‘Bottom Up Justification, Asymmetric Epistemic Push, and the Fragility of Higher Order Justification,’ *Episteme*, Online Early View, DOI: 10.1017/epi.2017.19.
- Kelly, T. (2010). ‘Peer Disagreement and Higher Order Evidence,’ in Feldman & Warfield (eds.), *Disagreement*. Oxford: Oxford University Press.
- Lasonen-Aarnio, M. (2014). ‘Higher-Order Evidence and the Limits of Defeat,’ *Philosophy and Phenomenological Research*, 88/2: 314-345.
- (forthcoming). ‘Enkrasia or Evidentialism?’, *Philosophical Studies*.
- Pollock, J. (1986). *Contemporary Theories of Knowledge*. Savage, MD: Rowman & Littlefield.
- Pryor, J. (2013). ‘Problems for Credulism,’ in Tucker (ed.), *Seemings and Justification*. Oxford: Oxford University Press.
- Skipper Rasmussen, M. (this volume). ‘Higher-Order Defeat and the Impossibility of Self-Misleading Evidence.’
- Titelbaum, M. (2015). ‘Rationality’s Fixed Point (or: In Defense of Right Reason),’ *Oxford Studies in Epistemology*, 5: 253-294.
- Worsnip, A. (2018). ‘The Conflict of Evidence and Coherence,’ *Philosophy and Phenomenological Research*, 96/1: 3-44.