

How your total evidence can mislead about itself

Alex Worsnip
UNC Chapel Hill

It's fairly uncontroversial that you can get misleading evidence about your evidence. Suppose your initial evidence does not support believing p , but then some expert tells you (in a rare lapse of judgment) that this initial evidence does support believing p . Then you have misleading "higher-order evidence" about your "first-order evidence" (and what it supports believing). That is:

Misleading Higher-Order Evidence.

- Your first-order evidence does not support believing p
- Your higher-evidence evidence supports believing that your first-order evidence supports believing p ¹

What is more controversial is whether such a phenomenon can ever result in a situation whereby both of the following are true:

Misleading Total Evidence (about Total Evidence).

- Your *total* evidence does not support believing p
- Your *total* evidence supports believing that your *total* evidence supports believing p ²

Let's say that the higher-order evidence is "decisive" with respect to believing a proposition just when it makes it the case that your *total* evidence supports believing that proposition. So, on our way of speaking, your first-order evidence can support believing a proposition without being decisive with respect to that proposition: this is so when it, on its own, supports believing a proposition, but the other evidence you have is such that overall, your total evidence does not support believing this proposition.

With this in mind, 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 decisive *both* with respect to believing the proposition that one's total evidence supports believing p (the "higher-order belief"), *and* with respect to believing the proposition p itself

This paper is greatly indebted to a conversation with Miriam Schoenfield, and an objection that she made to my (forthcoming). 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 David James Barnett and Jim Pryor for helpful discussion.

¹ Strictly speaking this is an *example* of misleading higher-order evidence, not a full definition of it. Here, one has misleading higher-order evidence that one's evidence supports believing p . But one might also have misleading higher-order evidence that one's evidence does not support believing p , misleading higher-order evidence that one's evidence supports suspending judgment about p , misleading higher-order evidence that one's evidence supports some particular credence in p , and so on. I'll continue throughout to focus on the example of misleading higher-order evidence that one's evidence supports believing p , but this is merely to fix ideas.

² The same remark made in fn. 1 applies, *mutatis mutandis*, here.

(the “first-order belief”).³ After all, “evidence of evidence is evidence.” So even when one lacks first-order evidence for *p*, the misleading higher-order evidence that one has might itself constitute decisive evidence for *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-decisive cases** – since they are cases where the misleading higher-order evidence is decisive *both* with respect to the higher-order belief *and* the first-order belief.

Secondly, one’s higher-order evidence might be decisive *neither* with respect to the proposition that one’s total evidence supports believing *p*, *nor* with respect to the proposition *p* itself. In such a situation, one’s higher-order evidence would itself support believing that one’s total evidence supports believing *p* – but this support would not be strong enough to be decisive with respect to this proposition. Remember that in a case of misleading higher-order evidence, 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 always 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-decisive cases** – since they are cases where the misleading higher-order evidence is decisive 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-decisive cases;⁴ others by claiming that all cases of misleading higher-order evidence are non-decisive cases.⁵ 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-decisive cases *or* non-decisive cases. That is: the misleading higher-order evidence *may* be decisive with respect to the higher-order belief, or may not be; the crucial claim is simply that *if it is* decisive with respect to the higher-order belief, it is also decisive 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-decisive cases**: cases where the misleading higher-order evidence is decisive with respect to the higher-order belief, but not the first-order belief. In my view, no-one has yet shown this decisively. I am going to try to do that, with a simple model.

³ It is a frequent source of confusion in the present debate that the adjectives “higher-order” and “first-order” can refer either to different bodies of *evidence*, or to different *beliefs* (or, if you prefer, propositions that are the objects of those beliefs). It can be easy to fall into the mistaken assumption that the higher-order evidence *just is* the evidence that bears on a higher-order belief, and the first-order evidence just is the evidence that bears on a first-order belief. (If that were so, (what I am calling) misleading higher-order evidence and misleading total evidence would be indistinguishable.) But that is not so. Higher-order evidence may bear upon a first-order belief (even if it must do so, in some sense, “indirectly”). Less obviously, but still plausibly, first-order evidence may bear upon a higher-order belief.

⁴ Cf., e.g., Feldman (2005), Bergmann (2005), and (more qualifiedly) Horowitz (2014).

⁵ Cf., e.g., Titelbaum (2015).

As a final preliminary, note 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.⁶ I concur here – evidence of evidence *is* (typically, *some*⁷) evidence – and my argument will not attempt to show otherwise. Indeed, it may even be that misleading higher-order evidence is *often* decisive with respect to the relevant first-order belief. What I aim to show is simply that misleading higher-order evidence can *sometimes* be decisive with respect to a higher-order belief, while not being decisive with respect to the corresponding first-order belief (that is, that it is sometimes singly-decisive). That is enough to show the misleading total evidence (about total evidence) is possible.

Let us work with a slightly more concrete case, adapted from Worsnip (forthcoming). In this case, Miss Marple and her niece Mabel visit a murder scene, where they survey the evidence. This (first-order) evidence does not support any particular verdict 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 evidence supports believing that the vicar did it.

Clearly, this case is one of misleading higher-order evidence. However, as the case was just stated, there seem to be few principled grounds for declaring that the case is either a doubly-decisive case, a non-decisive case, or a singly-decisive case – and it is only if it is a singly-decisive case that it is a case of misleading total evidence (about total evidence). However, I will give a simple model to show that, given a few eminently possible stipulations, Miss Marple’s testimony supports the higher-order belief (that the evidence supports believing that the vicar did it) *better* than it supports the first-order belief (that the vicar did it). I will then argue that this suggests that such a case can, in principle, be singly-decisive.

Three values are important here:

- X: The probability that the evidence supports believing that the vicar did it, conditional on Miss Marple saying that the evidence supports believing that the vicar did it
- Y: The probability that the vicar did it, conditional on the evidence supporting believing that the vicar did it
- Z: The probability that the vicar did it, conditional on the evidence not supporting believing that the vicar did it

Remember that what we are trying to show is that Miss Marple’s testimony (that the evidence supports believing that the vicar did it) supports the higher-order belief (that the evidence supports believing that the vicar did it) *better* than it supports the first-order belief (that the vicar did it). For the former level of support, we just look at the value of X. For the latter, we want the chance that Miss Marple is right that the evidence supports believing that the vicar did it *and* the vicar actually did do it, *plus* the

⁶ Cf., e.g., Kelly (2010); Pryor (2013: 99-100); Worsnip (forthcoming); Lasonen-Aarnio (forthcoming).

⁷ See Fitelson (2012) for a case where *pro tanto* evidence that there is evidence that p does not itself provide any evidence for p. But Fitelson explicitly says that this only works for *pro tanto* evidence. By his own lights, his case is not one where one has conclusive evidence that one’s evidence (conclusively) supports believing p. In fact, he says that he suspects that when one does have such conclusive evidence that one’s evidence supports believing p, *this fact is* always evidence for p (*ibid.*: fn. 2).

chance that Miss Marple is wrong that the evidence supports believing that the vicar did it *but* as it happens, the vicar did do it all the same. We can get this by multiplying X by Y, and multiplying (1-X) by Z, and then adding the two together – that is, by $(X.Y) + ((1-X).Z)$. So, in order for Miss Marple’s testimony to support the higher-order belief more strongly than it supports the first-order belief, we need values of X, Y and Z such that:

$$X > ((X.Y) + ((1-X).Z))$$

It’s easy to obtain values such that this is the case.⁸ For example, suppose X is 0.9, Y is 0.9 and Z is 0.5. That’s to say: given that Miss Marple says the evidence supports believing the vicar did it, there’s a 0.9 probability that the evidence does supports 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. Then, the value of the right-hand side is $(.9)(.9) + (.1)(.5) = .86$, which is less than the value on the left-hand side (0.9). Thus, given these values of X, Y and Z, the case is one where Miss Marple’s testimony supports the higher-order belief that the evidence supports believing that the vicar did it *better* than it supports the first-order belief that the vicar did do it. Assuming that there is some threshold of evidential support that is necessary for a belief to be supported on balance, this suggests that Miss Marple’s testimony could be decisive with respect to the higher-order belief, but not with respect to the first-order belief. For any reasonable threshold, there should be possible values of X, Y and Z that will do the job.

Note also that we can set possible (and, indeed, plausible) values for X, Y and Z that create a much bigger gulf between the support for the higher-order belief and the support for the first-order belief than that found above. In general, as X goes up, Y goes down, and Z goes down, the gulf will get bigger. For example, X might be 0.95, Y might be 0.8 and Z might be 0.2.⁹ Then, the value of the right-hand side is $(0.95)(.8) + (.05)(.2) = 0.77$. So we have now got a case where the probability of the higher-order proposition conditional on Miss Marple’s testimony is .99, and the probability of the first-order proposition conditional on her testimony is .77.

Two objections. First, one might point out that here we have only given the probability of the higher-order and first-order propositions conditional on Miss Marple’s testimony. But of course, this is not the only evidence that Mabel has. She also has the first-order clues gathered at the crime scene. And those also impact whether she meets the evidential threshold for belief. So, it might be said, I cannot conclude merely from the probability conditional on Miss Marple’s testimony meeting this threshold, that Mabel’s total evidence meets such a threshold, and thus that this higher-order evidence is decisive.

⁸ 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).

⁹ 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.

However, I did not in fact assume this. The present objection does nothing to impugn the result that the Miss Marple's testimony supports the higher-order belief (that the evidence supports believing that the vicar did it) *more strongly* than it supports the first-order belief (that the vicar did it). This opens the way for the possibility that this testimony pushes Mabel's evidential position over the threshold for believing that the evidence supports believing that the vicar did it, *without* pushing her evidential position over threshold for believing that the vicar did it. This is so even if whether she meets these thresholds is in part determined by other factors. The only thing that would block this argument would be if the antecedent evidential probabilities of both propositions (both their numerical values, and the resiliences of these values) were *guaranteed* to be such that Miss Marple's testimony, even though it supports the higher-order belief more strongly than the first-order belief, could not possibly change these probabilities such that the evidential threshold was met for the former, but not the latter. But the objector has given us no reason to make this strong, *ad hoc* supposition.

Here is the second objection.¹⁰ Recall that I am trying to show that one's total evidence could support believing that one's *total* evidence supports believing p, even though this is not the case. But Miss Marple's testimony is plausibly understood as being not about what the *total* evidence supports, but about what the *first-order* evidence supports. So, the objection continues, I have only given an example where one's total evidence supports believing that one's *first-order* evidence supports believing p, even though this is not the case. Call this misleading total evidence of first-order evidence. This is still, it bears stressing, more than mere misleading higher-order evidence (which was a case where one's *higher-order* evidence supports believing that one's first-order evidence supports believing p). But, it might be claimed, it still falls short of misleading total evidence *about total evidence*.

I admit that what the mathematical illustration above *immediately* yields is only misleading total evidence about first-order evidence. But, for any normal case, it is a short step from this to misleading total evidence about total evidence. Recall that in the case at hand, to have misleading total evidence about first-order evidence is to have misleading evidence that one's first-order evidence supports believing that the vicar did it. If the case were to provide an example of misleading total evidence about first-order evidence, but *not* misleading total evidence about total evidence, then, Miss Marple's testimony would have to be decisive with respect to the belief that Mabel's first-order evidence supports believing the vicar did it, but *not* be decisive with respect to the belief that Mabel's total evidence supports believing that the vicar did it. But how could this be so? Mabel's total evidence *just is*, we can stipulate, her first-order evidence *plus Miss Marple's testimony itself*. So the only way that her first-order evidence could support believing that the vicar did it, without her total evidence doing so, would be for Miss Marple's testimony – that the evidence supports believing that the vicar did it – to itself be countervailing evidence *against* the belief that the vicar did it. And so, the only way for her to be in a situation of misleading total evidence about first-order evidence, but not of misleading evidence about total evidence, would be for Mabel to have significant *evidence* that Miss Marple's testimony is itself evidence against the belief that the vicar did it.¹¹ Given how plainly Miss Marple's testimony is *not* evidence against such a belief, it is extremely difficult to see how this could be so. Perhaps a bizarre

¹⁰ Thanks to Miriam Schoenfield for this objection.

¹¹ What's more, to be of use to the objector, this evidence would have to achieve this *without* defeating the testimony's capacity to decisively support believing that the first-order evidence supports believing that the vicar *did* do it.

case where it is so could be cooked up¹² – but our present description of the Miss Marple case is not such a case.¹³ So this case, and indeed any case of misleading total evidence about first-order evidence which does *not* have this bizarre feature, will still be a case of misleading total evidence about total evidence. Thus, I conclude that cases of misleading total evidence about total evidence are possible.

In closing, let me briefly say a bit about why this matters.¹⁴ 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 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.
- Fitelson, B. (2012). ‘Evidence of evidence is not (necessarily) evidence,’ *Analysis*, 72/1: 85-88.
- Horowitz, S. (2014). ‘Epistemic Akrasia,’ *Noûs*, 48/4: 718-744.
- Kelly, T. (2010). ‘Peer Disagreement and Higher Order Evidence,’ in Feldman & Warfield (eds.), *Disagreement*. Oxford: Oxford University Press.
- Lasonen-Aarnio, M. (forthcoming). ‘Enkrasia or Evidentialism?’, *Philosophical Studies*.
- Pryor, J. (2013). ‘Problems for Credulism,’ in Tucker (ed.), *Seemings and Justification*. Oxford: Oxford University Press.
- Titelbaum, M. (2015). ‘Rationality’s Fixed Point (or: In Defense of Right Reason),’ *Oxford Studies in Epistemology*, 5: 253-294.

¹² Note, though, how odd it would be for someone who is trying to *limit* the capacity of evidence to mislead about itself to affirm this possibility.

¹³ By stipulation, in this case Mabel’s evidence is limited to her first-order evidence (the clues gathered at the crime scene) and Miss Marple’s testimony. Which one of these things could provide evidence that the testimony is evidence *against* believing that the vicar did it? Not the clues: they might themselves in some sense be evidence against believing that the vicar did it, or even for believing that Miss Marple’s testimony does not *decisively* support believing that the vicar did it, but they are not evidence that Miss Marple’s testimony is *itself* evidence *against* believing that the vicar did it. And not the testimony: the testimony that the evidence supports believing that the vicar did it is not evidence that *it itself* is evidence *against* believing that the vicar did it.

¹⁴ Here I’m very briefly summarizing the puzzle laid out in Worsnip (forthcoming) and in Lasonen-Aarnio (forthcoming).

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