

## Disagreement and Meta-Evidence\*

\*Draft version of a work in-progress. Do not use without permission.

Robert K. Whitaker

In the ongoing debate about rational peer disagreement, one of the central questions at issue is whether finding out what another person thinks about a proposition gives one any evidence relevant to that proposition. More specifically, does finding out that someone you take to be an *epistemic peer* disagrees with you about  $p$  give you any evidence about  $p$ ? Some have argued that such information need not significantly alter one's overall evidence assessment (e.g., Thomas Kelly), while others maintain that given certain plausible conditions, the views of others have defeating power for one's own beliefs (e.g., David Christensen, Adam Elga). A key player in this debate has offered a principle about evidence that has become central in this discussion, and it is this principle that I wish to examine here.

Richard Feldman was once an advocate of the so-called Conformist view—which holds that in cases of peer disagreement one should *conform* her view to that of her peer—but now thinks that there is no defensible overall solution to the problem of peer disagreement. He understands the epistemic significance of disagreement to be a matter of higher-order evidence that there is evidence against one's view. Discussing a case of peer disagreement involving “private evidence”—evidence that one cannot share with another—he says:

Each may have his or her own special insight or sense of obviousness. But each knows about the other's insight. Each knows that this insight has evidential force. And now I see no basis for either of them justifying his own belief simply because the one insight happens to occur inside of him. A point about evidence that plays a role here is this: evidence of evidence is evidence. More carefully, evidence that there is evidence for  $P$  is evidence for  $P$ . Knowing that the other has an insight provides each of them with evidence.<sup>1</sup>

This is Feldman's initial formulation of the now somewhat infamous principle “evidence of evidence is evidence.” I will label this principle the Meta-Evidence Principle (or EEE). In a couple of later papers, Feldman clarifies his intent in advancing this principle and responds to some objections.<sup>2</sup> The idea, he says, is not that one peer acquires the other's experiential evidence—“When you tell me that you have a headache, I don't thereby get your headache,” he says.<sup>3</sup> Rather, it is that your believing that  $p$  gives me reason to believe that there is some reason for  $p$ , which in turn gives me reason for  $p$ .

Several other commentators in the disagreement debate have had some (mostly negative) things to say about this principle. Much of this discussion centers on the relationship between evidence and justification. Plausibly, evidence typically *produces* justification for believing a proposition. If one is an evidentialist, like Feldman, then in fact evidence is *the only thing that can* produce justification. So on first glance, the principle seems

---

<sup>1</sup> Feldman, “Reasonable Religious Disagreements,” 208.

<sup>2</sup> Feldman, “Evidentialism, Higher-Order Evidence, and Disagreement,” 308-310. Also later paper...

<sup>3</sup> *Ibid.*, 309.

to imply that evidence that someone else has evidence for  $p$  gives one *justification* for the belief that  $p$ . But several commentators have wondered about this implication.

For example, Earl Conee notes that the principle does *not* entail justification because justification is a “summary evaluation” of a belief.<sup>4</sup> Meta-evidence could easily be defeated by other considerations. For example, if I learn that you, normally my epistemic peer about such things, believe that it will snow today (despite having been quite warm yesterday), then by the principle, I thereby gain some evidence that it will snow today. However, say I also find out that the weather app you’ve been using has a known malfunction that results in false predictions of snow. Then the evidence I gained from your evidence is canceled out, and so I am not justified in believing that it will snow.

Nonetheless, both Conee and Feldman suggest that evidence of evidence provides *some* defeasible evidence for a person—even if that evidence is *always* defeated. Conee considers a case where someone has some evidence that there is evidence for some proposition, but that person also knows that their belief in the proposition is a guess, with no better than a 50/50 chance of being true. This knowledge that it is a guess plausibly defeats the meta-evidence, but it nonetheless counts as evidence. (78) Similarly, Conee considers another case where one gains evidence that someone else has evidence for a proposition, in the context of being told about a hoax being perpetrated on that person. So Smith tells you that they made it seem to Jones as though Black was in the room today, when in fact they weren’t. By the meta-evidence principle, you now have evidence that Black was in the room—Jones’s evidence—but that evidence is immediately defeated by the fact that the source of your evidence—Smith’s testimony—is also defeating evidence for the proposition that Black was in the room. Feldman has a similar case: you read a retraction in the newspaper of a story from yesterday that you had not read. You now, says Feldman, get simultaneous evidence for and against what was reported yesterday, because the retraction gives you meta-evidence (and therefore evidence) for the content of yesterday’s report. But the retraction also gives you defeating evidence. As he says, “if you were asked whether you had any evidence at all in support of  $P$ , it is clear that the answer is ‘Yes’—you have evidence about yesterday’s newspaper report. To ignore this would be wrong.” (297) In other words, “defeated evidence is evidence.” (298) These cases also imply that evidence can be gained by *lying* (a point we will return to in a moment), but that such evidence is always or nearly always defeated by the fact that the person knows it is a lie. (88)

Another commentator, Alvin Goldman, interprets Feldman’s principle to mean that if one person “truthfully reports evidence  $Q$  concerning  $P$ , this evidence is also acquired by the hearer.”<sup>5</sup> (211) Now, this is not what Feldman means—and Goldman admits that it isn’t—but nonetheless Goldman makes a point that is relevant to Feldman’s intended usage: “Hearing such testimony may give the hearer default justification for believing  $Q$ , but such default justification can be defeated by other information in the hearer’s possession. *In that case,  $Q$  does not qualify as an item of evidence for the hearer.*” (italics mine) So according to

---

<sup>4</sup> Earl Conee, “Rational Disagreement Defended” (in Feldman & Warfield), 76-78, & Appendix 1, 84-89. Conee is discussing the principle:

EE: Evidence for Evidence. If  $S$  has evidence for the proposition that evidence exists in support of  $X$ , then  $S$  has evidence for  $X$  (76) Feldman notes this about his version of the principle as well...cite place...

<sup>5</sup> Alvin Goldman, “Epistemic Relativism and Reasonable Disagreement” (in Feldman & Warfield), 209-213

Goldman, apparently, evidence that is defeated—i.e. evidence that can play no justificatory role—is *not* evidence. We’ll return to this point below.

Jason Decker also briefly considers Feldman’s Meta-Evidence Principle.<sup>6</sup> Decker doesn’t say whether he thinks Feldman’s slogan is true, but he replies that meta-evidence *couldn’t* defeat one’s standing belief because one (ex hypothesi) has *no information* about the character of the other person’s evidence (as in the Righty and Lefty case). (771) Where one has no information, one ought to be agnostic, Decker argues, and agnostic in a “deep” sense, which Decker defines as assigning *no* credence to the proposition (by contrast, “shallow” agnosticism would be assigning a .5 credence to the proposition—i.e. suspending judgment because of *balanced* evidence). So Decker implies that meta-evidence, even if it exists, cannot play the defeating role that Conformists like Feldman would like it to play, because unless one knows what the evidence is, one is reasonable to remain deeply agnostic about it. This view, ironically, is itself agnostic between the two views on the Meta-evidence Principle just canvassed: Feldman’s (and Conee’s) and Goldman’s—i.e., Decker does not say whether he thinks that meta-evidence counts as defeated evidence, or is rather not evidence at all because it can play no justificatory role.

A more detailed critique of the Meta-Evidence Principle is provided by John Biro and Fabio Lampert.<sup>7</sup> They argue that when one learns that someone else disagrees with them about *p*, they only acquire evidence that that person *takes themselves* to have evidence about *p*. They do not thereby acquire evidence that *p* is *actually* true or false. (392) Biro and Lampert distinguish between two senses of evidence: “neutral evidence,” according to which someone *takes* something to tell in favor of *p* (whether or not it does), and “probative evidence,” according to which someone has evidence that *really does* tell in favor of *p*. (392) They then argue that Feldman’s principle only works if the second occurrence of “evidence” is interpreted in the probative sense: Evidence of *probative* evidence is evidence. (392) But of course we’re rarely (and in Feldman’s initial imagined cases, ex hypothesi *not*) in a position to know if the evidence the other person has is probative. And if it isn’t, then Biro and Lampert think we have no evidence relevant to what we should believe about *p* (the disputed proposition). So, as they say:

Imagine that we learn from a reliable source that a respected colleague has been claiming that LBJ was behind the assassination of JFK. We have evidence that he takes himself to have good evidence for the claim. Until we see what that evidence is, however, we have good reason to think him a whacko. Our evidence that he takes himself to have good evidence clearly does not mean that *we* have good evidence... (393)

All we have is evidence *that there is* (neutral) evidence pertaining to *p*, which should not, they think, alter our view of *p*.

---

<sup>6</sup> Jason Decker, “Disagreement, Evidence, and Agnosticism,” 771

<sup>7</sup> Biro & Lampert, “‘Peer Disagreement’ and Evidence of Evidence” They understand the principle to be: (EEE) Evidence that *S*’s near-peer disbelieves *p* is evidence against *p*. (391) But even “near-peer” is too strong. All that is required for the meta-evidence principle is that someone else (anyone) has some evidence. Presumably, that evidence is *\*stronger\** the more epistemically virtuous I take the other person to be, but if I can manufacture evidence via the principle by *\*lying\**, then anyone’s possession of evidence gives me evidence.

They clarify what they mean by “probative” evidence:

Two central facts about what we are calling probative evidence are, first, that no evidence can be such evidence for me if I do not know what it is and, second, that someone’s (even an expert’s) thinking that something is the case is not evidence that it is the case. (398)

This strikes me as a remarkably strong and highly counterintuitive claim. Here is a (skeptical) objection: Biro and Lampert’s view entails that probative evidence is impossible (or else reduces to something like perceptual or evidential seemings). This is because it is *always* possible that something we take to be evidence is not *really* evidence in the probative sense. So if probative evidence is required for defeat, then defeat will never happen, not even with those one takes to be epistemic superiors, because all that “superior” could mean for Biro and Lampert is “has evidence I lack.” But no one *can* have evidence I lack in the probative sense—or put differently, I can never know whether the evidence they think they have (or that I think they have) is probative. So I can never know if someone is my superior.

Hence, Biro and Lampert’s view of evidence entails an extreme dogmatism: it is never reasonable to alter one’s view because one never has evidence that her evidence is probative. To be sure, Biro and Lampert *think* that you can have evidence that your evidence is probative. They say, “to have evidence that there is probative evidence one has to know what the evidence supposedly probative is and to judge it probative.” (399) But if its being probative for me depends on my judging it to be so, then why can’t I judge the evidence of my peer’s report to be probative? How is this any different from me seeing a peer’s—or a superior’s—evidence myself and then judging it to be probative? In both cases, I could be wrong that it is probative, and if the person really is my peer (or my superior), then their report of the evidence should provide just as good (or better) a reason to believe that their evidence is probative, as if I had seen it myself. If it didn’t, then they would not really be a peer (or superior). So if I cannot know that my evidence is probative in the case of the report of a peer or superior, then I cannot ever know that my evidence is probative. And since Biro and Lampert think that if “one cannot judge whether [their evidence that p] is probative...one does not have evidence that p,” then one *never* has evidence for any proposition. (401)

Another way to put the point is that Biro and Lampert have a viciously circular notion of “epistemic superior.” As they say:

...we have no reason to believe that any evidence our disputant may have that we do not is probative evidence, something we assume with someone we believe, on good grounds, to be an expert. (402)

So we *assume* that if we believe someone to be an expert, then their reported evidence is probative. But how do we determine who is an expert? Well, by our judgment that their evidence is probative. This is what epistemic superiority *means* for Biro and Lampert: possession of more probative evidence than me. But *I* am the arbiter of what counts as probative, and I can never know that!

Another commentator, that I do not wish to dwell on here, argues that the Meta-Evidence Principle fails because it cannot pass a test of the transitivity of evidence. Branden

Fitelson argues that to refute the principle, we simply divide up the evidence into three parts—E1, E2, & E3, and then think of a case where E1 “supports” E2, and E2 “supports” E3, but E1 does not support E3, and voila, principle refuted.<sup>8</sup> I confess that this argument seems irrelevant to me—it apparently ignores the intended use of the principle in disagreement contexts and instead makes it about a failure of transitivity. But Feldman’s principle seems to me to have *nothing to do* with transitivity. He never claimed that evidential support relations were transitive. He claimed only that if a single person possesses meta-evidence about *p*, then that is the same as possessing evidence about *p*. The relation between E1 (the meta-evidence) and E3 (the resulting first-order evidence) here is one of *identity*, not transitivity. If E1 is a report, and E2 is another’s evidence that *p*, then E3 is just E1, with the qualification that the report is evidence for *p*—and not just for E2—and this regardless of what E2 is. E1 is evidence for *p*, not because it is entailed or supported by E2 (it might *not* be—see Biro and Lampert’s discussion of “probative” evidence), but simply because it was reported, and reports are evidential. So I fail to see the force of Fitelson’s objection. But other commentators that I respect, including Feldman himself, accept that Fitelson’s argument is relevant and successful at undermining at least one version of the Meta-Evidence Principle. So, on their testimony, I suppose I have evidence that there is something to be said for Fitelson’s argument, even though I don’t see it myself. (That’s a joke).

Let’s return to the point about automatically defeated evidence. In a couple of places, Feldman considers an objection to the Meta-Evidence Principle from Hud Hudson:

Suppose that at some APA conference we’re having our yearly beer together and I say, “Hey Rich, something kind of cool...it’s my birthday today!” And I do it in that winning and trustworthy way you’ve come to trust over the years. But I’m lying and I know I’m lying. I now have evidence for the proposition that evidence exists in support of the claim that it’s my birthday today, but I don’t have evidence (not even a little bit) for the claim that it’s my birthday today.<sup>9</sup>

Hudson’s point here is that if the Meta-Evidence Principle is correct, then Hudson *would* have evidence that it’s his own birthday, when he knows that it isn’t, which seems absurd.

Now, it seems to me that the way to respond to this objection is as follows: the Meta-Evidence Principle is not a norm governing one’s abilities to create evidence by lying. It is a principle about what the opinions of *others* give one reason to believe. This is due to the obvious fact (at least it seems obvious to me), which Hudson is here exploiting, that my *own* claims cannot create evidence *for me* for the content of those claims.

But this is not how Feldman responds to Hudson’s objection, and indeed it seems that he cannot. For if he were to admit that evidence is evidence only *for an agent* then the principle would fail to do the work he initially set for it, which is to show that evidence that a peer has evidence for *p* can be evidence for me that has defeating power for my own first-order evidence about *p*. Clearly if evidence for one person cannot be assumed to be evidence for another person, then this would not follow. So instead of making this (I think) sensible reply, he replies by biting the bullet: “I think that in his example Hud does get some

---

<sup>8</sup> Branden Fitelson, “Evidence of Evidence is not (Necessarily) Evidence,” *Analysis* 72, 1 (2012): 85-88.

<sup>9</sup> *Ibid.*

evidence for the proposition that it is his birthday today. But this is not problematic.”<sup>10</sup> It isn’t problematic, supposedly, because the evidence that Hudson receives through his lie is minimal and is easily swamped by the other evidence that he has that it is not his birthday. Feldman reminds the reader, rightly, that to have some evidence for *p* is not to be justified in believing *p*, since justification has to do with the balance of one’s total evidence. So, since Hudson does not obtain *justification* through his lie, there is no problem.

This seems to me to miss the force of the objection. The objection is not that Hudson obtains a *justified belief* that it is his birthday by lying about it being his birthday, but rather that he obtains *evidence* for the proposition that it is his birthday by lying about it being his birthday. Shifting attention away from this odd result and toward an even stranger result (justification) does not remove the original oddity. Surely evidence cannot be manufactured so easily. If it can, then all I need to do to obtain evidence for some view is to make a convincing show of evidence to someone else.

Actually, it’s even worse than this: the worry about justification may follow after all. Imagine that you are agnostic about some issue but would like to believe something about it for practical purposes (it would make Thanksgiving dinner conversation more palatable, let’s say). If you’ve examined the available evidence in the form of arguments and externally available reasons and concluded that suspense of judgment is called for, and if Feldman is correct, then you could move yourself away from agnosticism and toward your preferred belief state by simply lying to your family that you believe *p*. This is because they take you (let’s say) to be a reliable and honest reasoner, and so will take your report as evidence that there is a case to be made for *p*. So they have evidence that there is evidence for *p*; ergo, so do you. If the principle does not rule out gaining evidence through lying, and on the plausible assumption that reliably formed evidence for *p* produces (defeasible) justification for the belief that *p*, then you would be able to gain (admittedly low-level) justification for the belief that *p* by lying to your family. This is because your newly acquired evidence for *p* is undefeated and formed in a reliable way (i.e. lying in this way reliably produces evidence via the principle). Therefore, your belief that *p* is justified, or at any rate it is *more* justified than it was before lying, which suffices for my point.

Even if one were to insist here that there is some further distinction to be made between evidence and justification (which Feldman does), the central point stands: Feldman’s concession that evidence for *p* can be obtained in such a way is highly counterintuitive. But as I say, Feldman cannot respond differently, for if he concedes that Hudson’s claim creates evidence, but not evidence *for Hudson*, then he must admit that evidence of evidence is *not always* evidence *for* just anyone. And this would undercut the usefulness of the principle for his position, which is to establish that in a case of peer disagreement with private evidence, a peer’s reported disagreement can play a defeating role for one’s own private seemings.

It is also not clear what peer disagreement is actually evidence *of*. To see this, consider David Christensen’s mention of Feldman’s principle:

Suppose I have good reason to believe that my friend’s evidence, though different from mine, is just as good (for example, suppose my friend and I have done similar

---

<sup>10</sup> Ibid.

polls of distinct but comparable sample populations and have reached conflicting conclusions). Here, the reasons for our disagreement need not be cognitive--it could just be that one of the samples was not representative. But absent some special reason for treating her evidence or reasoning differently from mine, it seems clear that in evaluating explanations for our disagreement, I should regard our opinions as equally likely to be accurate, and thus I should alter my opinion toward hers.<sup>11</sup>

But this assumes that the peer's disagreement counts as evidence against one's own view about *p* (where *p* is a proposition about the sample population). Assuming that one has not seen the peer's evidence directly, why should this follow?<sup>12</sup> To see the point a bit more clearly, consider another case. Imagine that you are a defense attorney. Your client has claimed his innocence, and given the evidence you've seen, you're inclined to believe him. However, the prosecution is not behaving as if the defendant is innocent. In fact, they are behaving in a way that suggests they believe something you do not. You infer this because you know how *you* would behave as a prosecutor, given the evidence you have, and the prosecution is going well beyond that. So here you have some evidence in the form of the prosecution's seemingly odd behavior. But what exactly is this evidence *of*? The proposition in question here—the proposition in dispute between you and the prosecution—is: “The defendant is guilty.” Call that *p*. You deny *p*, while the prosecution affirms it. You each have evidence relevant to *p*, some shared and some, presumably, not (such as your own interviews with your client). Does the prosecution's odd behavior give you any evidence for or against *p* itself? I think not. It gives you, rather, evidence for another proposition: “The prosecution has information I lack.” Whether that information turns out to be evidence for *p* cannot be predicted before you know what it is. So it seems that here evidence that there is evidence for *p* (for someone else) turns out *not* to be evidence for *p* (for you). Evidence of evidence may be evidence, but not necessarily evidence *of* the disputed proposition(s), and not necessarily *for* both parties to a dispute.

But there is, it seems to me, something right about Feldman's EEE. I think what this is comes through nicely in a comment Peter van Inwagen makes about David Lewis:

Consider...the body of public evidence that I can appeal to in support of incompatibilism (arguments and other philosophical considerations that can be expressed in sentences or diagrams on a blackboard or other objects of intersubjective awareness). David Lewis “had” the same evidence (he had seen and he remembered and understood these objects) and was, nevertheless, a compatibilist. If I know, as I do, that David had these features (and this feature, too: he was a brilliant philosopher), that he had these features is itself evidence that is (or so it would seem to me) relevant to the truth of incompatibilism.<sup>13</sup>

---

<sup>11</sup> Christensen, “Epistemology of Disagreement: The Good News,” 211-212. He adds in a footnote here that “This is an example instance of Feldman's principle that ‘evidence of evidence is evidence.’”

<sup>12</sup> I take it that this is what Christensen intends. If one *had* seen the evidence directly, and found it to be “just as good” as one's own, then the uncontroversially correct response would be to suspend judgment, but not because of anything relating to the peer or her disagreement; it would simply be a case of inconclusive evidence.

<sup>13</sup> Van Inwagen, “We're Right. They're Wrong,” 26.

Notice that what is relevant to the justification of *p* (in this case, incompatibilism) is not disagreement itself, nor is it the fact that someone else possesses evidence for *p*. What is relevant to the justification of *p* is *that Lewis had these features*--namely, he had seen, remembered, and understood the evidence, and was brilliant. If none of these obtained, the question of the evidential value of his opinions would be moot. A being with these features creates an interesting epistemic problem, *even if that perspective is not actually instantiated*. So Feldman's principle is right in one sense: insofar as one is dealing with a perspective that is constituted by certain features of epistemic excellence, then the views of a person with that perspective may count as evidence for a disputed proposition. Note the implications here for Hudson's counterexample to EEE: if Feldman takes Hudson to have these features of epistemic excellence, then Hudson's testimony that it is his birthday does give Feldman evidence that it is Hudson's birthday. However, it does not give Hudson any evidence that it is his own birthday, because he knows himself to be lying, and therefore to be contravening the conditions for his own epistemic excellence, one of which is accurately representing the available evidence.

Perhaps this result is not surprising. Thomas Kelly gets at this point when he says: Let us set aside, for the moment, the special case of disagreement among peers, and reflect on a much more general question: in what circumstances does it make sense for me to treat the fact that someone else believes as she does as evidence for the truth of that which she believes? A true (although perhaps not especially informative) answer: exactly when I take her belief to be a reliable indication of how things stand in the relevant part of reality. Thus, suppose that I know, on the basis of extensive past experience, that, when my weather forecaster judges that it will rain the next day, it tends to rain 80 percent of the time. In that case, I will treat her judgments to the effect that it will rain as evidence that it will rain, inasmuch as I take there to be a positive correlation between the two. Notice that, in this respect, there is absolutely nothing special about the way in which the judgments of another person come to count as evidence."<sup>14</sup>

Again, it is (roughly) the reliability of the person that makes their opinion evidentially valuable, and this in turn depends on the specific features of their reasoning habits. So we might modify Feldman's principle to yield a more plausible version as follows:

MEEE: For peers S1 and S2, evidence of S2's trustworthy judgment that *p* is evidence for S1 that *p*.

This takes into account the perspectival nature of evidence (that it's *for* S1, though not for S2); it specifies the object of the evidence (the proposition *p*); and it emphasizes that in order to count as evidence for someone else, one's evidence must be rooted in a trustworthy judgment, which implies various features of epistemic excellence.

---

<sup>14</sup> Kelly, "Peer Disagreement and Higher-Order Evidence," 132. If Kelly is correct here, this is further reason the disagreement problem does not reduce to a problem of testimony.