RICHARD FELDMAN

EVIDENTIALISM, HIGHER-ORDER EVIDENCE, AND DISAGREEMENT

ABSTRACT

Evidentialism is the thesis that a person is justified in believing a proposition iff the person's evidence on balance supports that proposition. In discussing epistemological issues associated with disagreements among epistemic peers, some philosophers have endorsed principles that seem to run contrary to evidentialism, specifying how one should revise one's beliefs in light of disagreement. In this paper, I examine the connection between evidentialism and these principles. I argue that the puzzles about disagreement provide no reason to abandon evidentialism and that there are no true general principles about justified responses to disagreement other than the general evidentialist principle. I then argue that the puzzles about disagreement are primarily puzzles about the evidential impact of higher-order evidence – evidence about the significance or existence of ordinary, or first-order, evidence. I conclude by arguing that such higher-order evidence can often have a profound effect on the justification of first-order beliefs.

Evidentialism is a widely accepted thesis about epistemic justification. A simple way to formulate it is as follows:

1. Evidentialism: S is justified in believing P at t if and only if S's evidence at t on balance supports P.

A more complete statement of the view adds that other attitudes – disbelief and suspension of judgment – are justified when the evidence goes against the proposition or is neutral.[1] An additional detail that will matter in what follows is that the degree of justification a person has for believing a proposition is determined by how strongly the evidence supports that proposition.[2]

The evidentialist view about disagreements is quite simple: people are justified in believing what their evidence supports when they are involved in a disagreement. Consider a person who believes P and learns that someone relevantly like herself (an epistemic peer) disbelieves P. The evidentialist view is that the person is justified in continuing to believe P if and only if the evidence the person has after learning of

 DOI: 10.3366/E1742360009000720

https://doi.org/10.3366/E1742360009000720 Published online by Cambridge University Press

the disagreement on balance supports P. The person is justified in ceasing to believe P if and only if the evidence the person has after learning of the disagreement on balance does not support P.[3]

Much of the recent debate among epistemologists about disagreements has been about the merits of various principles specifying the epistemic status of a person's belief after learning about peer disagreement. These principles say that if certain conditions hold, then one is, or one is not, justified in continuing to hold one's belief. Some of the principles seem to run contrary to evidentialism, and in the first section of this paper, I will examine the connection between evidentialism and these principles. I will defend the view that there are no true special principles about justified responses to disagreement. That is, there are no general and widely applicable principles specifying what attitudes are justified for someone who encounters a disagreeing peer. The relevant truths are facts about what some evidence supports, the general evidentialist principle itself, and the implications of these factors. I will argue that this failure to find general principles about justified responses to disagreements should come as no surprise – something similar is true in other domains.

While there may be no special general principles about justified responses to disagreements, there are facts about the evidential impact of disagreements. I remain sympathetic to the conciliatory view that evidence of peer disagreement is often significant evidence against one's view.[4] Often, this evidence is higher-order evidence – evidence about the significance of one's first-order evidence. In the later sections of this paper, I will discuss several issues associated with this kind of evidence.

1. EVIDENTIALISM AND PRINCIPLES ABOUT DISAGREEMENT

As an example of the kind of principle about which I am skeptical, consider

2. If S justifiably believes P immediately prior to t, and S learns at t that an epistemic peer with respect to P who shares S's evidence concerning P disbelieves P, then S is not justified in believing P at t (then S is justified in suspending judgment about P at t).[5]

Philosophers who are inclined to accept conciliatory views – views that imply that learning of peer disagreement undermines justification – typically endorse more qualified views than (2). I will consider some revisions below, but it will be useful to begin with (2).

If people "share evidence" only if they have exactly the same evidence, then it is likely then there are no actual cases in which the condition in the antecedent of (2) is satisfied. The spirit of the proposal might be captured in a principle that weakens this condition somewhat, perhaps to something like this:

https://doi.org/10.3366/E1742360009000720 Published online by Cambridge University Press

3. If S justifiably believes P immediately prior to t, and S learns at t that an epistemic peer with respect to P who has evidence comparable to S's concerning P disbelieves P, then S is not justified in believing P at t (then S is justified in suspending judgment about P at t).

The idea of "comparable evidence" is obviously vague, but it can plausibly be interpreted to cover cases that are prominent in the recent literature. For example, if two philosophers examine all the same arguments concerning some issue and come to different conclusions about that issue, we might have a case in which two people have comparable evidence. The example in Peter van Inwagen (1996) of his dispute with David Lewis about free action is as good a case as any. There are, of course, differences between van Inwagen and Lewis in this case: their responses to the arguments differ. Presumably, this difference does not imply that their evidence is not comparable.[6]

(3) may also apply to other prominent cases in which people do not actually share evidence with others, but instead have evidence that is in certain respects similar to the evidence of their peer. One such example, in David Christensen (2007), involves two people who compute their shares of a restaurant bill. They have the same information about the total amount of the bill and the number of people among whom it must be divided. They arrive at their conclusions about each person's share by doing arithmetic calculations of the same sort. We might say, then, that they have "comparable evidence." The same is true of a case in which people surveying the same scene have different perceptual experiences, and thus different beliefs.[7] They have the same kind of evidence in the same kind of circumstances, and we might regard this as comparable.

One might take (2) and (3) to be anti-evidentialist principles. On this view, the idea is that regardless of the merits of one's overall evidence subsequent to learning about the disagreement, one loses one's justification for one's initial belief. This would be a fact about epistemic justification not fully determined by one's evidence. However, the principles need not be taken in this way. Instead, one might think that they reflect the implications of evidentialism when applied to cases of the sort described in the antecedents of the principles. On this view, (2) and (3) follow from evidentialism combined with certain principles about evidential support. The principles would specify the evidential force of the information about peer disagreement. Specifically, an evidentialist who endorses (2) and (3) is committed to principles such as

4. If S justifiably believes P immediately prior to t, and S learns at t that an epistemic peer with respect to P who shares S's evidence concerning P disbelieves P, then S's evidence at t on balance does not support P (then S's evidence at t is counterbalanced with respect to P).

https://doi.org/10.3366/E1742360009000720 Published online by Cambridge University Press

5. If S justifiably believes P immediately prior to t, and S learns at t that an epistemic peer with respect to P who has comparable evidence concerning P disbelieves P, then S's evidence at t on balance does not support P (then S's evidence at t is counterbalanced with respect to P).

In my view, there is nothing in the facts about disagreements that warrants abandoning evidentialism. I will not argue for this here, except to note the oddity of a view according to which in general one is justified in following one's evidence but that learning about peer disagreement warrants abandoning one's evidence. I know of nothing to support this kind of exceptionalism. If this is right, evidentialists who accept principles like (2) and (3) are stuck with principles like (4) and (5).

On reflection, (4) and (5) strike me as highly implausible. For them to be true, the evidence obtained when one learns about peer disagreement must overwhelm anything else that might be true about one's evidence in order to assure that one's overall evidence fails to support one's initial belief. A variety of examples suggest that this is not the case. Here are a few:

- Suppose that a person learns that an epistemic peer disagrees on an issue, but at the same time learns that this is the lone dissenter among a huge group of peers. Add to the story, if it helps, that those who agree on this particular issue often disagree among themselves about other issues, so it is not a case of people who simply take their cues from one leader. The huge preponderance of evidence about peer beliefs supports the person's initial belief. It is hard to see why the one person who disagrees should be treated as deferentially as the principles under discussion suggest. Indeed, if respect for one's peers is supposed to be so valuable, then (4) and (5) – hence (2) and (3) – seem ill-advised because they attach so little epistemic import to the beliefs of peers who agree in this case.
- Suppose that a person is well-justified in believing a proposition and learns that a peer with similar evidence disagrees. However, the disagreeing peer also says that he is hardly confident of his own conclusion and can easily understand someone disagreeing with him. While this might provide some slight reason for the person to become a little less confident of his belief, it is hard to see why this kind of imbalanced situation makes suspension of belief the justified attitude. That is, belief, though perhaps less confident belief, seems to be the justified attitude. Again, if respect for the attitudes of peers is the consideration driving conciliatory views, then principles such as those under consideration are not the ones to adopt.
- Suppose that a person learns that a peer with the same (or extremely similar) evidence disagrees with him about an issue, but also learns something that undermines the credibility of the peer in this particular instance. Assuming that "peer", as it occurs in (4) and (5), is used in a way that allows for the possibility of cases in which a peer is known to be performing below par, this kind of case is easy to imagine.

https://doi.org/10.3366/E1742360009000720 Published online by Cambridge University Press

All these cases turn on a common theme: What's justified depends upon one's overall evidence, and this larger body of evidence can offset whatever influence the evidence of peer disagreement might have.

There are several ways one might revise (2) or (3) in an effort to get around these problems. These attempts will yield principles that fall into two main categories: those that strengthen the antecedent either by adding more conditions that must apply to a disagreement for it to require revision of the sort specified in the consequent or by modifying the definition of "peer", and those that weaken the consequent by changing the epistemic outcome disagreement brings about. It's also possible to revise the principles to say that if the conditions in the antecedent obtain, then those specified in the consequent "usually" obtain or will obtain "other things being equal." I will address each of these possibilities.

Before turning to these revisions of (2) and (3), it will be useful to take note of a truth that may have led people, myself included, to assert things along the lines of (2) and (3): learning of peer disagreement provides a reason to modify one's view and often has a profound impact on what one's overall evidence supports. This can be formulated as follows:

6. If S believes P, and S learns at t that an epistemic peer with respect to P who shares S's evidence concerning P disbelieves P, then S acquires some evidence against P.

Of course, the fact that S believes P doesn't really matter in (6). It is simply learning of a peer who disbelieves P that provides evidence for him against P. And the underlying idea here can be put not as a principle about justification at all but instead as a principle about epistemic support, such as

7. The proposition that S's peer who shares S's evidence concerning P disbelieves P is evidence against P.[8]

I believe that (7) expresses the key evidential fact about disagreements and that the attempted repairs of (2) and (3) that I will examine next fall into one or another of several categories; they are either false, or simple restatements of (7) and its implications, or so highly limited in application that they fail to spell out the truth behind conciliatory views. Given (7), there is no need for these principles.

One way to strengthen the antecedents of (2) and (3) is to add more descriptive facts about the situation. For example, one might add facts that rule out the elements that were present in the three examples I described above. Thus, the conditions would specify that the subject does not also learn that a huge preponderance of peers share his own view, or that the peer who disagrees is far less confident of his judgment than the subject is of her judgment, or that the disagreeing peer is somehow impaired or less than trustworthy on this occasion. It should come as no surprise that no finite list of factors will suffice to yield true principles. There will always be the possibility of some other factor being present

https://doi.org/10.3366/E1742360009000720 Published online by Cambridge University Press

that provides the subject with evidence that she ought not defer to the peer this time. For example, a successful principle must rule out the possibility that the subject also learns that epistemic superiors all share her view.

One might hope to identify some way to characterize all these factual situations that might obtain and that would provide evidence that counteracts the evidence provided by the disagreeing peer. The history of attempts to formulate principles of this general sort does not provide grounds for optimism here. What might seem to do the trick is a principle simply stating that the subject who learns of the disagreement does not also get any additional evidence supporting her original belief. The result is a principle such as

8. If S justifiably believes P immediately prior to t, and S learns at t that an epistemic peer with respect to P who has evidence comparable to S's concerning P disbelieves P, and S does not learn anything that counteracts the evidence provided by this, then S is not justified in believing P at t (then S is justified in suspending judgment about P at t).

An alternative way to formulate this might be in terms of defeaters, or the lack of them:

8a. If S justifiably believes P immediately prior to t, and S learns at t that an epistemic peer with respect to P who shares S's evidence concerning P disbelieves P, and the evidence provided by this fact is not defeated, then S is not justified in believing P at t.

These principles are not true if the evidence provided by peer disagreement can ever be weaker than the subject's initial evidence supporting her belief. If the subject begins with good reason to believe P, perhaps supported by peer agreement, then the new evidence provided by the peer disagreement might weaken (or partially defeat) her evidence without making it the case that her belief is no longer justified (nor make it that she should suspend judgment). It need not be that powerful.

This last consideration suggests that we might find a true principle by weakening the consequent of (8) so that the principle states only that the subject becomes less justified in believing P when the antecedent conditions obtain:

8b. If S justifiably believes P immediately prior to t, and S learns at t that an epistemic peer with respect to P who has evidence comparable to S's concerning P disbelieves P, and S does not learn anything that counteracts the evidence provided by this, then S becomes less justified in believing P at t.

A similar principle could also be formulated in terms of lack of defeat. I think that (8b) may well be true.[9] But if it is, this is because it is a direct consequence of

https://doi.org/10.3366/E1742360009000720 Published online by Cambridge University Press

evidentialism and (7) (the fact that peer disagreement provides evidence). Consider any proposition anyone justifiably believes. If the person acquires some evidence against that proposition and the person learns nothing that counteracts this new evidence, then the person's overall justification is reduced. (9) is hardly some special insight into the epistemic impact of disagreement. It is simply a special case of a general truth about evidence combined with (7).

Another way to try to repair principles such as (2) and (3) is to add qualifiers specifying that it is "usually" or "typically" the case that when the antecedent is true, then the consequent is true. Such principles are sufficiently vague to make decisive evaluation extremely difficult. But what seems clear to me is that the resulting principles are just vague and indefinite ways to say what (7) says. That is, the principles amount to roundabout ways of saying that evidence of disagreement is evidence against what one believes. Nothing beyond (7) is advanced by such principles.

Another way to try to repair (2) or (3) is to modify the definition of "peer". This amounts to strengthening the antecedent, and discussion of the resulting principles will go over essentially the same ground as the discussion just concluded. Almost any way one strengthens the definition – adding conditions on peerhood that rule out flaws or defects that undermine credibility – will be inadequate. What will work is an account of "peer" that renders the principle true by definition. That is, if one simply defined "peer" as a person such that evidence of disagreement with that person in the specified circumstances renders one's belief not justified, then the principle will be true. But this is a just a matter of defining one's way to a true principle. I don't believe that this is the sort of substantive principle epistemologists seek.

A few additional points about strengthened definitions of "peer" merit attention. I believe that the concept of an epistemic peer is a vague concept of which we do have reasonably good grasp. The key element of any account will be a specification of how much alike two individuals have to be in order to count as peers. If the standard of similarity is set very high – requiring exactly the same evidence, abilities, epistemic virtues, and the like – it will turn out that there are no actual peers. In that case, the resulting principle will apply to few or no actual cases. Perhaps the resulting principle will be true, but it will hardly capture the point of the conciliatory view due to its having such limited application. If the standards for being a peer are lower, then the principles will be subject to the kinds of objections considered above. The key point being that other evidence also obtained may counteract the evidence obtained from the peer.

There is, of course, no precise notion of what an epistemic peer is. The only reason to try to make this concept more precise is in an effort to make principles such as those under discussion immune to criticism. But that project is futile. Moreover, in any realistic case, when we encounter someone who is more or less like us, disagreement will have some evidential impact, and how much impact it will have will depend upon the details of the situation. What really matters is how

https://doi.org/10.3366/E1742360009000720 Published online by Cambridge University Press

credible the person is on this topic on this occasion (or, more precisely, what it is reasonable for the subject to believe about this). But the opinions of people who are not our peers can have impact as well. This point supports the general idea I'm advocating here: it is the total evidence that matters. Restricting attention to unusual cases involving full-fledged peers, where this is given an interpretation making few people peers, yields principles that might be true, but will have little application and will fail to capture the key idea of conciliatory views.

Somewhat similar points apply to efforts to defend the principles by manipulating the evidential similarity condition in the principles. As noted early on, restricting ourselves to principles about cases in which people have exactly the same evidence will make for principles that do not apply to any real cases. If, for example, my evidence includes anything along the lines of "It seems to me that ____ " or "It appears to me that ____ ", then, since you will not have exactly my experience, we will not have the same evidence. Once the principles are expanded to include similarity of evidence, all the complexity of evidential situations comes into play again, and no principle short of the general evidentialist principle is going to be correct.

Trying to be clearer about the concepts of an epistemic peer, sharing evidence, or having comparable evidence can be useful exercises. Of course, these are technical concepts and any definitions will be to some extent stipulative. I think that the only reason to try to make these concepts more precise is in an effort to make principles like (2) and (3) work. You have to come up with definitions that avoid the sorts of examples I characterized above. However, I believe that (6) and (7) are fine as they stand, using intuitive (and imprecise) interpretations of these terms.

These considerations lead me to doubt that any principle along the lines of (2) and (3) will be both true and widely applicable and not a simple direct consequence of evidentialism. My doubt about the truth of any interesting and informative principles arises because I think that for any set of conditions you can specify for the antecedent of the principle, there will be some possible way to add to the evidence specified in that antecedent some additional evidence that will yield a different status than the one specified in the consequent. There is always the possibility of something defeating or counteracting the evidence specified in the antecedent. There may be some true principles with highly detailed antecedents that apply only to a narrow range of special cases. Other principles will simply amount to restatements of evidentialism conjoined with (7).

I do not contend that I have shown that there are no true general and widely-applicable "epistemic status" principles specifying the epistemological implications of learning about peer disagreement. I have merely cast doubt on a few principles and raised a question about whether there is any way to repair those principles. I believe, however, that some modest support for this "no principles" view can be gained by examining some other situations and by considering an analogy. I will begin with the analogy.

https://doi.org/10.3366/E1742360009000720 Published online by Cambridge University Press

Consider a principle about the moral status of promise-keeping such as

9. If S promises to do A, then S ought to do A.

From a broadly consequentialist perspective, this principle, taken as a universal generalization, will appear implausible. It may be that most promises ought to be kept, but this is true only because promise keeping tends to have good consequences. Or, if one attaches value to promise-keeping itself, one will likely think that the positive value of promise-keeping can be outweighed in some cases by other considerations. In other words, one might think that being a case of promise-keeping is something that counts in favor of an action, but it is not by itself decisive. This might yield a moral principle about promises analogous to (7). Principles such as (9) that state generalizations about the moral status of all cases of promise-keeping are implausible because they fail to take into account the potential value of other factors. This is analogous to what I think goes wrong with the epistemic status principles about disagreement that I have considered.

In my view, epistemic status principles in other domains are also doomed to failure. There is unlikely to be much dispute about highly specific principles such as

10. If S learns that the *New York Times* reports that P, then S is justified in believing P.

I suspect that no one will accept (10). Even if we assume that the *New York Times* is a reliable source that merits our trust, (10) is not true. There are many situations in which one will have evidence that outweighs whatever support the report provides. What is true in this case is, perhaps, something like

11. The proposition that the *New York Times* reports that P supports the proposition that P.

Of course, one might argue that it is only a contingent fact that the *Times* is a reliable source. If that's the case, then one might doubt that (11) is true. Instead, the nearby truth might be something along the lines of

12. The conjunction of the proposition that the *New York Times* reports that P and the proposition that the *New York Times* is reliable supports the proposition that P.

Just as there are not likely to be true, general, and applicable special evidential status principles about propositions one finds to be reported in the newspaper, there are no true evidential status principles about propositions one learns to be subject to peer disagreement. Instead, reliable newspaper reports, as well as information about peer disagreement, have evidential impact, though what the overall result of learning such information will be depends upon the rest of one's evidence.

https://doi.org/10.3366/E1742360009000720 Published online by Cambridge University Press

I believe that essentially the same points apply to principles governing what one might take to be "fundamental" sources of evidence. Consider:

13. If at t S seems to remember that P, then S is justified in believing P at t.

Again, it is not controversial that this is false. The evidence provided by apparent memory can be outweighed by other evidence. Strengthening the antecedent in various ways will not yield truths. For example, restricting the topic of the remembered proposition will not suffice, nor will adding anything about the strength or clarity of the apparent memory. Any such evidence can, in principle, be overridden. What is true is simply this:

14. Seeming to remember that P provides evidence for P.

There is, I think, no true general and widely applicable epistemic status principle about memory. The nearby relevant truth is one about epistemic support.

One might think that in the case of memory, and perhaps in the case of the *New York Times*, there are acceptable epistemic status principles that add a "no-defeaters" condition to the antecedent. Consider:

15. If at t S seems to remember that P, and the evidence provided by this fact is not defeated by any other evidence S has, then S is justified in believing that P.

According to one standard interpretation of "defeated", this is true. On this interpretation, all it means to say that some other evidence, E′, defeats the support E provided to P is that (E and E′) does not support P. Given (14), the first conjunct of the antecedent of (15) implies that S has some evidence supporting P. Given the interpretation of "defeat" now under consideration, the second conjunct implies that nothing in the rest of S's evidence, when conjoined with E, fails to support P. (This allows that S does have some evidence against P.) It follows that S's evidence on balance supports P. And given the weak (i.e., not "knowledge level") notion of justification intended in (15), it follows that its consequent is true. Thus, under this interpretation, (15) simply follows from (14) and the definition of "defeat". It is really just an instantiation of a general principle stating that if a person has some evidence supporting a proposition, and that evidence is not defeated (in the sense specified), then the person is justified in believing the proposition. Again, the key fact is the evidential support fact, and principles such as (15) are, if true, simple consequences of these facts. This mirrors the situation in the case of disagreements.

Consideration of these cases provides some modest support for the idea advanced here, that there are no true general and widely applicable principles governing epistemic disagreements. There is, of course, considerable vagueness in this claim. However, I hope that the idea is reasonably clear: there are no simple facts about the epistemological significance of evidence of peer disagreement except for the fact (and its implications) captured by (7) that evidence that a peer disagrees has evidential impact.

https://doi.org/10.3366/E1742360009000720 Published online by Cambridge University Press

In defending these views about disagreement, I am denying principles sometimes defended under titles such as "Split the Difference" or "Equal Weight". However, I do not think that I am disagreeing with a key thought that motivates such views. I am struck by the existence of disagreement and puzzled about its epistemic import. In thinking through this topic, a key question is whether it is reasonable to maintain one's beliefs in the light of peer disagreement. While I am not endorsing universal principles asserting that it is never reasonable to maintain one's belief, I am arguing that evidence of peer disagreement is evidence against one's original belief. It is consistent with this that, in many cases, it is strong evidence against one's original belief, strong enough to render that belief no longer justified. Thus, I continue to endorse the view that peer disagreement is epistemically significant and threatens to undermine the justification of many beliefs in areas of controversy.

2. DISAGREEMENTS AS A HIGHER-ORDER EVIDENCE PROBLEM

In this section I will discuss the way higher-order evidence enters into the issues surrounding disagreements. Higher-order evidence can be characterized as follows:

16. Higher-order evidence: Evidence about the existence, merits, or significance of a body of evidence.

A clear case of higher-order evidence is the evidence one obtains when one learns how someone else assesses a body of evidence – what that person thinks it supports. A second example would be evidence that someone has unspecified evidence supporting some proposition. For example, if someone tells you, "Jones has some evidence supporting P, but I don't know what it is," then you have some evidence about the existence of evidence for P. A third kind of case of higher-order evidence is evidence that another person has a certain sort of evidence, you do know what it is, but you do not have that evidence yourself. Consider a case in which someone looking at an object reports, "It looks blue to me." You do not have any visual evidence that the object is blue. But you have some evidence about the existence of the other person's visual evidence.

It is worth noting that when you get first-order evidence, you will almost always get higher-order evidence as well, at least if you have a certain level of sophistication. Suppose you see an object that looks blue. Assume the blue look is evidence that the object is blue. Then, if you are sophisticated, you have learned that someone (you) has evidence for the proposition that the object is blue. And note also that awareness of this fact does not provide you with even more reason to believe that the object is blue. (Assume for the sake of discussion that the mere blue look does provide evidence that the object is blue. This is a contentious matter that I will not address here.)

https://doi.org/10.3366/E1742360009000720 Published online by Cambridge University Press

It is clear that some of the controversy about disagreements turns on points about the epistemological significance of higher-order evidence. To think this through, it will be useful to distinguish some importantly different kinds of examples. One way to categorize the cases is as follows:

17. Evidential significance disagreements: disagreements in which the participants disagree about the significance or evidential impact of a body of evidence.
18. Unspecified evidence disagreements: disagreements in which one learns that another person has some unspecified evidence supporting a belief competing with one's own belief.
19. Specified evidence disagreements: disagreements in which one learns that another person has some specific evidence supporting a belief competing with one's own belief (but does not acquire the other person's evidence).

Evidential significance disagreements introduce a kind of puzzle that the latter two kinds of disagreements do not. The latter two are really just about how to weigh various bits of competing evidence. For example, suppose we are looking at the same object and report having different beliefs about its properties: it does not look blue to me, but you say that it looks blue to you. Here, I do not get your experiential evidence, but I get reason to believe that you have evidence of a certain kind. The question then arises about what influence this newly acquired evidence should have on my belief that the object is blue. This is a perfectly good question, with no simple answer. Some of the relevant considerations are easy enough to identify. What I know (or have reason to believe) about our relative competence on this sort of topic, our current conditions, how careful we've been, how likely you are to be kidding, and so on will affect the final calculation. As I argued (or asserted) in Section 1, there is no simple principle specifying how all these factors will add up. Essentially the same thing is true if I learn that you have some unspecified evidence that has led you to believe something contrary to what I believe. This also presents an issue about how to weigh competing bits of evidence on a topic. This is a familiar, though difficult, issue.

Evidential significance disagreements introduce a different sort of issue. These cases are more puzzling than the others, and the second-order aspect of the evidence is more significant. The key fact here is that the first-order evidence may actually support a proposition, and the disagreement provides evidence that the first-order evidence does not support that proposition. The question that then arises is what to make of the total evidence in such cases. I will return to this topic in Section 3.

I think that some of the disputes in the literature may be about the weight of higher-order evidence – some thinking evidence from others weighs more heavily than others think it does. Although it is hard to prove, I suspect that those who have said things that seem to support Equal Weight or Split the Difference views are inclined to think that, at least in some significant range of cases, this higher-order

https://doi.org/10.3366/E1742360009000720 Published online by Cambridge University Press

evidence is quite weighty. Dissenters tend to deny this. But I should note that sometimes in listening to discussions of the topic, it seems that some people attack the "Equal Weight" view, denying that you always have to count another person's view as equally weighty as your own (thereby forcing suspension of judgment in certain kinds of cases) and others (from the other side) attack the "no difference" view, arguing that you do have to give those views some weight. There need not be any disagreement among advocates of those two positions, although there could be a difference in their views about how things are to be weighed.

I think that one's higher-order evidence often makes a great deal of difference for the status of one's first-order beliefs. I don't know that there is much dispute about the proposition that such evidence makes some difference – only those who say that you can always just ignore information like this from others really do disagree with this. One reason for consensus on this is that there is no doubt that the kind of information one gets in these examples has epistemic significance when it is not part of a disagreement. That is, learning what a peer thinks about evidential significance when one does not have a prior opinion does provide some evidence about the topic. But if the evidence about the other person would be epistemically significant in cases in which one did not have a view of one's own, it is hard to see why it has no impact when one does. You have to think that somehow having your own view renders it completely insignificant. That seems implausible.

In any case, my point here is just that the disagreement issue is often about the significance of higher-order evidence, and that there are these very different kinds of cases that should be distinguished.

3. ON THE EVIDENTIAL SIGNIFICANCE OF HIGHER-ORDER EVIDENCE IN DISAGREEMENTS

In this section and the one that follows, I will address a few issues concerning the evidential significance of higher-order evidence of the sort described in Section 2. This will not come close to being a comprehensive analysis of the topic. I will begin with issues associated with evidential significance disagreements.

It is worth noting that other epistemological problems turn on similar issues. The most notable such problem is skepticism in some guises. One thing that can happen when studying skepticism is that one can become aware of considerations that seem to call into question the relationship between one's evidence and the conclusions one is inclined to draw on the basis of that evidence. In some cases, skeptical arguments appear to present people with reasons to doubt that the evidence they have for some class of propositions actually does provide justifying evidence for those propositions. Thus, the arguments can seem to provide people with reasons to doubt that their perceptual evidence really does support ordinary external world propositions, that the evidence in familiar inductive inferences really does support the conclusions routinely drawn, and so on. One question that arises in these circumstances is what impact this higher-order evidence about the

https://doi.org/10.3366/E1742360009000720 Published online by Cambridge University Press

merits of one's first-order evidence has on the epistemic status of one's first-order beliefs.

One way to approach the issues raised by these considerations is by introducing the concept of "levels incongruity." We can characterize this as follows:

20. Levels Incongruity: believing P while believing that your evidence does not support P, or believing P while considering the proposition that your evidence supports P and suspending judgment on that proposition.

The key issue about levels incongruity is whether these combinations of attitudes can be justified. The issue is brought into sharpest relief by consideration of cases in which one's first-order evidence actually supports P, one (initially, at least) believes P, and then one gets what would appear to be good reasons to doubt that the first-order evidence supports P. Examples are cases in which a valued and trusted peer makes a convincing case that your evidence does not support what you think it does The remarks that follow concern only cases of this sort. For the reasons outlined in Section 1, I doubt that there will be any simple general principles that determine the justified responses in these cases. But this does not preclude making any relatively general comments.

There are, I believe, three views one might have about any such case:

i) One is justified in maintaining one's belief that P and justified in believing (accurately) that the first-order evidence does support P (in spite of the apparent reasons to the contrary).
ii) One is justified in maintaining one's belief that P and also justified in suspending judgment or in disbelieving that the first-order evidence supports P.
iii) One is not justified in maintaining one's belief that P and is justified in disbelieving or suspending judgment on the proposition that the first-order evidence supports P.

The only way I see to defend (i) in a particular case is by arguing that one is somehow automatically justified in believing the truth about the epistemic support relation in question. One might hold that (in certain cases, at least) these relations are "transparent" or *a priori* and that this implies that no evidence to the contrary can render the truths unjustified. This strikes me as highly implausible, though I do not have any decisive arguments against it. It seems clear that propositions about epistemic support relations, just like propositions about any other subject matter, can be supported in any number of ways, including testimony, reasonable but ultimately unsound arguments, and the like. (i) assumes a kind of impossibility to be led astray, a power that we just don't have.

Solution (ii) introduces a kind of Moorean puzzle. While there is no logical inconsistency in the attitudes one has in (ii), it is odd to maintain attitudes that have the kind of levels incongruity this implies. This view seems to have the implication that reflection on epistemic support relations has no impact on the justification of one's first-order beliefs in these cases. Imagine this attitude applied to another

https://doi.org/10.3366/E1742360009000720 Published online by Cambridge University Press

kind of case, one in which a person does not encounter disagreement but instead is simply reflecting alone. The person reasonably comes to the conclusion that his evidence does not support P. And he then stops believing P. Approach (ii), applied to this case, would have the proper analysis be that this person was reasonable in drawing his conclusion about his evidence, but not reasonable in forming a first-order belief in light of that evidence. This isolation of levels seems a clear mistake to me. There would be something extremely odd about criticizing a person who justifiably believes that his evidence does not support a proposition for failing to believe that proposition.

This leaves response (iii). In general, this strikes me as a plausible response in those cases in which the evidence does make disbelieving or suspending judgment on the evidential support proposition reasonable. Note that to defend (iii) as the right response to these cases does not imply that it is reasonable to disbelieve or suspend judgment simply because a peer asserts that one's own assessment of the evidence is incorrect. It may be that whatever evidence comes from one's own reflection on the topic often carries greater evidential weight.

This does not amount to arguing that the original evidence should be ignored. The idea is not that one has to weigh the original evidence supporting the proposition against the higher-order evidence that goes against the proposition. Rather, the higher-order evidence serves as a defeater of the original evidence, in the sense that the conjunction of the original evidence and the higher-order evidence fails to support the proposition. It is, in some ways, like what happens to the belief that an object is red when one learns that a red light is shining it.[10]

Thus, I conclude that evidence that one is in an evidential significance disagreement with a peer does provide significant evidence that can undermine a belief that is in fact supported by one's original first-order evidence.

4. WEIGHING EVIDENCE

I turn finally to cases in which one gets specified or unspecified evidence that one's peer has evidence against one's own conclusion. As I said earlier, these cases raise familiar but difficult questions about how to weigh competing pieces of evidence. I have nothing noteworthy to say about this general issue here. I simply want to respond to one line of thought about the reasoning I've endorsed on this topic and to consider one interesting kind of example.

In Feldman (2006, 223) I claimed that when one learns that another person has evidence supporting a proposition, one has evidence supporting that proposition oneself. I used "Evidence of evidence is evidence" as a slogan to capture this idea. Perhaps the slogan is somewhat misleading. Here's a way to make this more precise:

21. Evidence of evidence is evidence: If S has evidence for the proposition that evidence exists in support of P, then S has evidence for P.

https://doi.org/10.3366/E1742360009000720 Published online by Cambridge University Press

This does not mean that if I learn that you have evidence for P, I thereby obtain your evidence. If there is experiential evidence, then when you have a headache, you have experiential evidence supporting the proposition that you have a headache. When you tell me that you have a headache, I don't thereby get your headache. But I do then have reason to think that you have a headache. I get evidence, but not your evidence.

Hud Hudson has raised an objection.[11] He writes:

> 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.

His point is that he knows that I have evidence for the proposition that it is his birthday, but he himself doesn't have evidence for that proposition. So this may look like a counterexample to (21). I believe that it is not. I think that in his example Hud does get some evidence for the proposition that it is his birthday today. But this is not problematic. Given everything he learns at the time, his evidence does not on balance support that conclusion any more than it did previously. There are several principles, all false, that might lead one to think otherwise. Here are two:

22. If S has some evidence for P, then S is completely justified in believing P.
23. If S has some evidence for P, then S is slightly justified in believing P.

(22) is obviously false, given the existence of competing evidence. (23) is equally obviously false, given that one can have evidence that is massively swamped by evidence for the denial of what the initial evidence supports. The true principles in the area are these:

24. If there is something, Q, that S knows, and Q supports P, then S has some evidence for P.
25. If S has some reason to think that someone has a reason to believe P, then S has some reason to believe P.

These principles make it relatively easy to have evidence for something. In fact, each of us may have some evidence for almost everything. But I think that this is true. Many such reasons are easily and quickly and completely defeated – that is, there is something else, D, you know that combined with it is not a reason. So, E is a reason, E&D is not. But the presence of D does not make E not a reason. It makes it a defeated reason. The "is a reason for relation" is timeless and eternal and necessary. S has a reason for P provided there is something S "has", and it's a reason for P.

To sustain this line of thought, there are lots of conversational puzzles we have to avoid. To note one: if I have a defeated reason and you ask if I have a reason

https://doi.org/10.3366/E1742360009000720 Published online by Cambridge University Press

for P, I'd mislead if I simply said "yes". But I would not have spoken falsely. In fact, I can say that I have reason, but it is defeated. If you asked if there were considerations in favor of P, I'd mislead or worse if I just said "no."

Another point is worth noting. Acquiring evidence for a proposition does not always increase one's level of justification for that proposition. For example, suppose I learn that you know that I believe P. This is evidence for both P and the proposition that I believe P. However, by acquiring this evidence I do not become better justified in believing either P or the proposition that I do believe P. Evidence does not add up in a simplistic way.[12]

Thus, I stand by the "evidence of evidence is evidence" principle, properly understood. But this tells us nothing about the bottom line on the overall significance of the kind of evidence one has in cases of disagreement. The fact is that this is complicated. The evidence from peers with whom one disagrees does count. But the main thing to say is that it should be added to the mix. This may seem disappointing – you want some crisper principles that you can object to. But, as I said, we shouldn't expect them.

Finally, I want to consider another kind of example that raises an important point about unspecified evidence disagreements. Ernest Sosa points out that in many

> cases of deep, important controversy . . . . the evidence on which we base belief in our side of a controversy need only be inscrutable, for whatever reason, or at least sufficiently hard to uncover. For it cannot then be displayed for reflection on how well it supports the content of our belief. (forthcoming)

He goes on to say that our evidence is "subtle and complex" and that

> Our basis for believing as we do on such questions generally fails to be fully formed and operative in one fell swoop. Light here dawns gradually. A belief forms in us over time through the subtle influence of diverse sources. Some are testimonial, others perceptual, others inferential, and so on. The belief might owe importantly to the believer's upbringing, or to later influence by his community.

Because of this, we cannot fully share our evidence in these cases. Sosa is almost surely right about this. But it is a mistake to conclude that this somehow protects the justification of our beliefs in these cases.

Suppose that Sosa is right about our inability to report our evidence. I therefore know that you have been unable to tell me your whole story. I also do not know my own entire evidential story. But I also know your strengths. I have evidence that you have evidence for your conclusion. I know that there are two lives, one that led to belief in P, one that led to belief in ~P. Both are, I have agreed, generally epistemically worthy lives. I am more familiar with one – my own. I do not see how this makes me better justified in believing what the path I have followed has led to than what the path that you have followed has led to.

https://doi.org/10.3366/E1742360009000720 Published online by Cambridge University Press

5. CONCLUSION

In this paper I have attempted to defend a conciliatory view about disagreements without defending any general, widely-applicable principle about justified responses to disagreements. Instead, I have tried to provide some support for the view that the key truth behind the conciliatory view is captured in the claim that the proposition that a peer disagrees with you concerning a belief is evidence against the proposition you believe. This fact, combined with evidentialism, implies that when you learn of a disagreement, you acquire evidence against what you believe. Exactly what impact this will have on what is supported by your overall evidence is a complex matter and will depend upon the details of each specific case. I have also argued that the evidence acquired in cases of disagreement is higher-order evidence. I have attempted to distinguish a few kinds of higher-order evidence and to develop a few points about how higher-order evidence impacts the justification of lower-level beliefs.[13]

REFERENCES

Christensen, David. 2007. "Epistemology of Disagreement: The Good News." *Philosophical Review* 116: 187–217.

Conee, Earl and Richard Feldman. 2004. *Evidentialism: Essays in Epistemology*. Oxford: Oxford University Press.

Elga, Adam. 2007. "Reflection and Disagreement." *Noûs* 41: 487–502.

Feldman, Richard. 2006. "Epistemological Puzzles about Disagreement." In S. Hetherington (ed.), *Epistemology Futures*, pp. 216–36. New York: Oxford University Press.

Feldman, Richard. 2007. "Reasonable Religious Disagreements." In L. Antony (ed.), *Philosophers Without Gods*, pp. 194–214. New York: Oxford University Press.

Feldman, Richard. In preparation. "Evidence of Evidence is Evidence."

Kelly, Thomas. Forthcoming. "Peer Disagreement and Higher-Order Evidence." In R. Feldman and T. Warfield (eds.), *Disagreement*. Oxford: Oxford University Press.

Matheson, Jonathan. 2009. "Conciliatory Views of Disagreement and Higher Order Evidence." *Episteme, A Journal of Social Epistemology* 6(3): 269–79.

Sosa, Ernest. Forthcoming. "The Epistemology of Disagreement." In *Armchair Philosophy*. Princeton: Princeton University Press.

van Inwagen, Peter. 1996. "It is Wrong, Always, Everywhere, and for Anyone, to Believe Anything, Upon Insufficient Evidence." In J. Jordan and D. Howard-Snyder (eds.), *Faith, Freedom, and Rationality*, pp. 137–54. Lanham, MD.: Rowman & Littlefield.

NOTES

1 I will formulate this discussion in terms of belief, disbelief, and suspension of judgment. I believe that almost everything I discuss could also be formulated in terms of degrees of belief.

https://doi.org/10.3366/E1742360009000720 Published online by Cambridge University Press

2 See Conee and Feldman (2004) for a more complete discussion of evidentialism.
3 Others have endorsed the same view. For explicit discussion of this, see Kelly (forthcoming).
4 I come close to endorsing a conciliatory principle in Feldman (2006, 2007). Others who have defended similar views include Christensen (2007) and Elga (2007).
5 The version of the principle making suspension of judgment the justified attitude has no plausibility whatsoever if S's initial belief is not justified. If a person believes P but is justified in disbelieving P, it surely is not the case that learning of peer disagreement would make suspension of judgment justified. If anything, learning of the disagreement would make disbelief even more strongly justified. However, if one starts with a justified belief, and then learns of disagreement, then "splitting the difference" might make suspension of judgment the justified attitude.
6 Principle (3) applies to both van Inwagen and Lewis only if both are initially justified in their beliefs. Some may doubt that this is true. A similar point applies to the examples in the next paragraph. To accommodate this, one might explore variants of (3) that omit "justifiably" from the antecedent.
7 My own "dean on the quad" case in Feldman (2007) illustrates the idea.
8 I express evidential support relations simply as relations between propositions. What is justified for a person depends upon what evidence the person has.
9 There are two reasons to doubt that (9) is exactly right. For one thing, it's possible that acquiring some undefeated evidence against P does not always weaken one's position. Reading multiple copies of a newspaper may provide "new" evidence for or against various propositions, but eventually the reports have no additional epistemic impact. For another thing, it is possible that the person already has something that counteracts the new evidence. Perhaps the principle formulated in terms of "lack of defeat" handles this second case better than (9) does.
10 I discuss this sort of point in Feldman (2007). For a more thorough discussion of the topic, see Matheson (this issue).
11 Via email.
12 I expand upon these points in Feldman (in preparation).
13 Earlier versions of this paper were presented at the *Episteme* Conference at Northwestern University in June 2009, at the Workshop on Disagreement at the University of Copenhagen in August 2009, and at a conference on Responsible Belief in the Face of Disagreement at the VU University, Amsterdam, Netherlands in August 2009. Rik Peels commented at the latter conference. I am grateful to him and to the participants of all three conferences for extremely helpful comments.

Richard Feldman is Professor of Philosophy and Dean of the College at the University of Rochester. His research interests are in epistemology. His publications include *Reason and Argument*, 2nd Edition (1999), *Epistemology* (2003), and *Evidentialism* (with Earl Conee) (2004).

https://doi.org/10.3366/E1742360009000720 Published online by Cambridge University Press