

Disagreement and Higher-Order Evidence

Yan Chen & Alex Worsnip

Forthcoming in Maria Baghramian, J. Adam Carter & R.A. Rowland (eds.),
The Routledge Handbook of The Philosophy of Disagreement

Penultimate version – please cite published version if possible

In the contemporary epistemological literature, peer disagreement is often taken to be an instance of a more general phenomenon of “higher-order evidence.” Correspondingly, its epistemic significance is often thought to turn on the epistemic significance of higher-order evidence in general. This chapter attempts to evaluate this claim, and in doing so to clarify some points of unclarity in the current literature – both about what it is for evidence to be “higher-order,” and about the relationship between disagreement and higher-order evidence. We will begin by considering some candidate definitions of “higher-order evidence,” and offering our own definition that attempts to capture the phenomenon of interest (§1). We will then consider, in light of this definition, whether disagreement and its epistemic significance are best-understood as a kind of higher-order evidence (§2). We’ll argue that although peer disagreement can be epistemically significant qua higher-order evidence, this role doesn’t *exhaust* its epistemic significance, and that it can also serve as straightforward first-order evidence. Finally, we’ll suggest in §3 that inattention to this latter point has made broadly conciliatory views about peer disagreement seem somewhat easier to resist than they in fact are.

1. What is higher-order evidence?

Despite a ballooning literature on higher-order evidence, the term ‘higher-order evidence’ is surprisingly hard to define precisely, and is, in practice, used in a confusing variety of different ways. To make things worse, the variety in the different ways that it is used is rarely acknowledged. So we will try to provide some clarity about what might be meant by this term.

We can begin with an easier notion, namely that of a higher-order *proposition*. Suppose we have a proposition (*q*) that states a claim about the evidential, justificatory, or rational status of [belief in] some other proposition (*p*) [that is not itself about the evidential, justificatory, or rational status of some still further proposition]. Then, *p* is a first-order proposition and *q* is a higher-order proposition. For example, suppose we have the first-order proposition *that a Democrat will win the 2024 US Presidential election* – call that proposition *FOP-D* for short.¹ Then, the following (among others) are all higher-order propositions:

- Your evidence supports [believing] *FOP-D*.
- Your evidence does not support [believing] *FOP-D*.

Thanks to Adam Carter for comments on a previous draft.

¹ In labelling propositions throughout, we’ll preface first-order propositions with ‘FOP’ and higher-order propositions with ‘HOP’.

- It would be justified/rational for you to believe *FOP-D*.²
- It would not be justified/rational for you to believe *FOP-D*.
- Your existing belief in *FOP-D* is justified/rational.
- Your existing belief in *FOP-D* is unjustified/irrational.

Now, what about higher-order *evidence*?³ One popular way to characterize higher-order evidence is as evidence *about* whether the kinds of higher-order propositions listed above are true.³ But it is not very easy to say what makes a piece of evidence *about* whether such propositions are true. One could construe ‘about’ very narrowly so that the evidence in question has to literally have one of these higher-order propositions as part of its content. So, for example, if a reliable expert tells you *that your evidence supports [believing] FOP-D*, that is clearly evidence about whether your evidence supports [believing] *FOP-D*. But this seems much too restrictive to capture the full phenomenon that philosophers have meant to pick out with the term ‘higher-order evidence’. For example, getting evidence that you formed a belief under the influence of a reason-distorting drug is supposed to be a paradigm instance of higher-order evidence, but it doesn’t have any claim about whether your belief is justified or rational *as part of its content*; rather, it just constitutes *evidence* that your belief may not be justified or rational. That suggests a much broader, and appealingly simple, precisification of the definition of higher-order evidence:

HOE₁. Some item of evidence E is higher-order evidence iff E is evidence for or against some higher-order proposition.

However, this now seems unhelpfully broad as characterization of higher-order evidence. Indeed, it may even be that on this definition, all evidence is higher-order evidence. For example, suppose a new opinion poll comes out showing a large lead for Democrats over Republicans in the 2024 US presidential election race. This is a textbook example of the sort of thing that epistemologists consider first-order evidence. And indeed, if we define first-order evidence in the same way that HOE₁ defines higher-order evidence, this evidence will count as first-order evidence, since the poll is evidence for *FOP-D*, a first-order proposition. But notice that the poll is plausibly also (at least some) evidence *that your evidence supports FOP-D*. (If this isn’t obvious, consider that after learning about the poll result, you should be more confident that your evidence supports *FOP-D*.) Thus, the poll also fits the definition of higher-order evidence given by HOE₁.⁴ And similarly for pretty much any ordinary piece of evidence. Epistemologists have typically wanted to pick out a more distinctive phenomenon with the term ‘higher-order evidence’.

A different characterization of higher-order evidence would be as follows:

² Throughout, we use ‘rational’ and ‘irrational’ in their “substantive” sense, whereby they are synonymous with “justified” and “unjustified”.

³ See, e.g., Skipper & Steglich-Petersen (2019: 1).

⁴ Dorst & Hedden (forthcoming) make a similar point. They conclude that (almost) all evidence is higher-order evidence. Here they seem to presuppose a definition of higher-order evidence along the lines of HOE₁. But one could equally take this point to show that we need a different definition of higher-order evidence to zero in on the phenomenon of interest.

HOE₂. Some item of evidence E is higher-order evidence iff E is evidence for or against some higher-order proposition *without* constituting evidence for or against the first-order proposition that the higher-order proposition refers to.

There is some evidence that plausibly fits this characterization. Consider the following example:

Wishful Thinking. You strongly believe (i.e., have a very high credence) that a Democrat will win the 2024 US election, on the basis of what *seems* to you to be very strong evidence that a Democrat will win the 2024 US election. However, your friend then confronts you with evidence (perhaps, past track-record data) that shows that you are highly prone to engage in (subconscious) wishful thinking about political outcomes, where you overestimate the probative force of evidence for the outcome you hope for, and underestimate the probative force of evidence against it, leading you to be more confident that your hoped outcome will occur than the evidence warrants you in being. And indeed, in this case, you strongly hope that a Democrat wins the 2024 US election.

In this case, it seems that the evidence your friend confronts you with is at least *some* evidence for the following higher-order proposition:

HOP-D. My strong belief (very high credence) that a Democrat will win the 2024 US Presidential election is irrational.

After all, the fact that you have engaged in wishful thinking so many times in the past, and that this is exactly the sort of proposition that you are inclined to engage in wishful thinking about (viz. one predicting a political outcome that you hope will occur) is at least *some* evidence that you have engaged in wishful thinking this time too – in which case, you are more confident than you should be, and hence, your (strong) belief is irrational. However, the fact that you are prone to engage in wishful thinking is *not*, it seems, evidence either for or against the first-order proposition *FOP-D*, namely that a Democrat will win the 2024 US election. (Nevertheless, it is thought, it may have some rational bearing on your belief in *FOP-D*; more on this soon.) Hence it plausibly fits the characterization of higher-order evidence given by HOE₂.

Some epistemologists have given a characterization of higher-order evidence close to HOE₂.⁵ However, whereas HOE₁ was too broad to zero in on the phenomenon of interest, HOE₂ arguably seems too narrow. To see this, consider the following case:

Expert Detective. Sherlock Holmes is an expert detective who is incredibly reliable in determining what the evidence supports when it comes to all matters to do with crime. You and Sherlock visit a murder scene together. After reviewing all the evidence together, Sherlock tells

⁵ Christensen (2010) and Staffel (2021) both define higher-order evidence as evidence about the agent's own rationality or reasoning, implicitly excluding evidence about what the agent's evidence supports. The former kinds of evidence will typically satisfy HOE₂, while the latter typically won't.

you, “the evidence you have just seen strongly supports the hypothesis that the priest committed the murder.”

Notice that in this case, what Sherlock tells you is not that the priest *did* commit the murder, but rather *that your evidence supports* the hypothesis that the priest committed the murder. (If it turned out that the priest did not commit the murder, but the evidence you had seen *misleadingly* strongly suggested that it did, then Sherlock would have been speaking truly, not falsely.) Because of this feature, many epistemologists cite cases like **Expert Detective** as cases of higher-order evidence.⁶ But note that Sherlock’s testimony plausibly does *not* fit the definition given by HOE₂. Sherlock’s testimony is evidence for the following proposition:

HOP-M. Your evidence supports (believing) that the priest committed the murder.

After all, Sherlock simply asserted *HOP-M*, and by stipulation, he is incredibly reliable about such matters. However, as well as constituting evidence for *HOP-M*, Sherlock’s testimony is *also* plausibly at least some evidence for the following proposition:

FOP-M. The priest committed the murder.

As we’ve said, Sherlock’s testimony is evidence for *HOP-M*, the proposition that your evidence supports (believing) that the priest committed the murder. But, when the evidence supports (believing) *p*, more often than not, *p* is actually true. So Sherlock’s reliably telling you that your evidence (albeit, in a sense, indirectly) supports (believing) that the priest committed the murder is plausibly at least some indication that the priest actually did commit the murder; that is, it is evidence for *FOP-M*. Consequently, Sherlock’s testimony is not higher-order evidence according to HOE₂.

A broader issue with both definitions that we’ve considered so far is that they don’t capture the way that a single piece of evidence can be higher-order evidence *with respect to one proposition* while being first-order evidence *with respect to another*. For example, Sherlock’s testimony is higher-order evidence with respect to the first-order proposition *FOP-M*, but there seems to be a good sense in which it is first-order evidence with respect to the higher-order proposition *HOP-M*. After all, Sherlock’s testimony simply asserts that *HOP-M* is true. Since Sherlock is reliable about matters like *HOP-M*, his asserting it is evidence for *HOP-M* in the same way that reliable testimony about *any* kind of proposition is evidence for that proposition.⁷ It just happens to concern a higher-order proposition; it plays no distinctively higher-order role with respect to that proposition.

To fix the forgoing problems, we suggest the following characterization of higher-order evidence:

⁶ Cf., e.g., Horowitz (2014: 728-9); Worsnip (2018: 24-25).

⁷ For our justification for treating ordinary testimony as first-order evidence, see footnote 22.

HOE₃. Some item of evidence *E* is higher-order evidence with respect to a proposition *p* iff (i) *E* is evidence for or against a higher-order proposition concerning *p*; and (ii) *E*'s (putative) rational bearing on doxastic attitudes toward *p* (if any⁸) is explanatorily derivative on (i).

When we talk about a 'rational bearing' here, we mean to refer to the capacity of the evidence in question to change (or contribute toward changing) the evidential, justificatory, or rational status of (doxastic attitudes toward) *p*. We use this umbrella term to accommodate both the possibility that higher-order evidence does this by actually constituting evidence for or against *p* (as is plausible in Expert Detective), or that it does so in some other way (as is plausible in Wishful Thinking). It can thus capture both cases as cases of higher-order evidence. Let us explain this in a bit more detail.

Consider first Expert Detective. Here, the explanation of why Sherlock's testimony is evidence for *FOP-M* (i.e., the proposition that the priest committed the murder) goes through the fact that it is (regular, first-order) evidence for *HOP-M* (i.e., the proposition that your evidence supports believing that the priest committed the murder). Roughly, it runs as follows: since Sherlock is extremely reliable about matters concerning what the evidence supports, his saying that *HOP-M* is (regular, first-order evidence) for *HOP-M*. But what *HOP-M* says is just that your evidence supports *FOP-M*. Thus, since the (putative) fact that your evidence supports *FOP-M* is itself a fairly reliable indicator that *FOP-M*, Sherlock's testimony is also some evidence for *FOP-M*. Thus, the rational bearing it has on *FOP-M* is that of constituting evidence for *FOP-M*. Nevertheless, the fact that it constitutes evidence for *FOP-M* is, as we just saw, explanatorily derivative on the fact that it constitutes evidence for *HOP-M*. Thus, it fits the definition given by HOE₃.

Now consider Wishful Thinking. Here, the fact that you are prone to wishful thinking is (regular, first-order) evidence for *HOP-D* (i.e., the proposition that your belief that a Democrat will win the 2024 US Presidential election is irrational). As we've already noted, the fact that you are prone to wishful thinking is *not* evidence for or against *FOP-D* (i.e., the proposition that a Democrat will win the 2024 US Presidential election). However, many (though not all⁹) epistemologists hold that it nevertheless has some other kind of rational bearing on your belief in *FOP-D*.¹⁰

There are different ways of trying to spell this out. One idea is that the information that you are prone to wishful thinking constitutes an *undercutting defeater* (either in whole or in part) for whatever justification you previously had for believing *FOP-D*.¹¹ On this diagnosis, it is like discovering that a wall that appears red is being illuminated by a red light. While this isn't itself evidence that the wall is *not* red – its being illuminated by a red light tells us nothing about whether it is red underneath – it nevertheless changes the rational status of believing that the wall is red by *undercutting* the justification

⁸ This qualification allows someone who denies the epistemic significance of higher-order evidence (cf. footnote 9) to still call it higher-order evidence. However, one could also omit the qualification and say that those who deny the epistemic significance of higher-order evidence don't really think that there is any higher-order evidence as such.

⁹ Cf. Lasonen-Aarnio (2014), Titelbaum (2015), and Weatherson (2019: esp. ch. 8).

¹⁰ Does holding that it has a rational bearing on your belief in *FOP-D* despite not being evidence for or against *FOP-D* violate a plausible version of evidentialism, on which the rationality on your belief in *FOP-D* is a matter of whether it is supported by your total evidence? Not necessarily, for it might *change what your total evidence* supports without itself being evidence for or against *FOP-D*. This is most obvious on the undercutting defeat proposal noted below: there, the idea is that an undercutter makes it the case that your *other* evidence (to the extent that it previously supported *p* does not do so).

¹¹ See e.g. Avnur & Scott-Kakures (2015). For the general notion of undercutting defeat see Pollock (1986: 39).

you had previously: in the knowledge that the wall is being illuminated by a red light, its appearing red is no longer strong evidence that it is actually red, and hence it is no longer rational to believe that the wall is red. Others hold that higher-order evidence of this sort cannot be assimilated to the phenomenon of undercutting defeat,¹² and that another diagnosis is needed: among the proposals are that it involves a distinct kind of “dispossessing” defeat,¹³ that (in cases like Wishful Thinking) it requires you to *bracket* your assessment of the first-order evidence,¹⁴ or that it constitutes a special, non-evidential reason to suspend judgment.¹⁵

We cannot adjudicate between these different proposals in the space available here. The important point for our purposes is that on all of them, the discovery that you are prone to wishful thinking has a rational bearing on (your belief in) *FOP-D* in virtue of the fact that it suggests that your belief in *FOP-D* is (or may well be) irrational: that is, in virtue of its being evidence for *HOP-D*. For example, on the “bracketing” proposal, it is precisely because the discovery that you are prone to wishful thinking is evidence that you have *overestimated* the strength of the (first-order) evidence for *FOP-D* – thus ending up with irrationally high confidence in *FOP-D* – that you are required to bracket your assessment of that first-order evidence. Similarly for the other proposals. That is enough for us to see that the Wishful Thinking case can also be accommodated as higher-order evidence on the definition given by HOE_3 .

A final complication bears stressing because it will be important shortly in discussing disagreement. We intend the definition of higher-order evidence given by HOE_3 to allow for the possibility that sometimes, an item of evidence might be both higher-order *and* first-order evidence with respect to the very same proposition p . This would be so if it has some rational bearing on (doxastic attitudes toward) p derivatively on its status as evidence for a higher-order proposition concerning p , **and** it also has some separate, independent rational bearing on (doxastic attitudes toward) p that is not derivative on its status as evidence for a higher-order proposition concerning p . (We’ll later suggest that peer disagreement is an instance of this.) So we don’t mean clause (ii) of HOE_3 to be read as requiring that *all* of E’s rational bearing on p must be derivative on its status as evidence for the higher-order proposition in question. But it will be true to say that all of its rational bearing on p qua higher-order evidence is derivative on its status as evidence for the higher-order proposition in question.

2. Is disagreement first-order evidence or higher-order evidence (or both)?

¹² Cf. Feldman (2005: 111-3); Christensen (2010: 193-5); White (2010: 585); Lasonen-Aarnio (2014: 317-8); González De Prado (2020: 327-8); Weatherson (2019: 193); Lord & Sylvan (2021: 124). Briefly: the difference is supposed to be that undercutting defeaters actually “sever” support relations between evidence and doxastic attitudes that would otherwise be present. By contrast, higher-order evidence like that received in the Wishful Thinking case suggests that there was *never* such a support relation. It indicates that one’s belief was *never* rational to start with (which is not so in the Red Light case). Moreover, slightly more controversially, if the higher-order evidence is misleading, and one’s belief *was* supported by one’s first-order evidence, arguably the higher-order evidence relation doesn’t make that cease to be the case.

¹³ González De Prado (2020).

¹⁴ Christensen (2010).

¹⁵ Lord & Sylvan (2021).

Having given a definition of higher-order evidence, we now move on to address the issue of whether the epistemic significance of disagreement should be assimilated to that of higher-order evidence.

2.1 *The current orthodoxy*

As anyone reading this handbook will know by now, there has been a great deal of attention to the epistemic significance of disagreement, addressing how we should adjust our beliefs when we find ourselves in disagreement with others, especially with our epistemic peers. The current views in the literature fall on a spectrum. At one end, the most extreme kind of “Conciliationist view” argues that disagreement always requires one to conciliate with one’s peer, by suspending one’s belief in p or significantly downgrading one’s credence in p .¹⁶ At the other end, the most extreme kind of “Steadfast view” argues that disagreement is never, in itself, a reason to adjust one’s beliefs or credence.¹⁷ Between these extremes, various other views are possible, on which peer disagreement requires significant conciliation under some circumstances but not others.¹⁸

Despite their divergence, the majority of the existing proponents of these various views seem to agree that *to the extent* that disagreement is epistemically significant, it is significance a kind of higher-order evidence that, in one way or another, bears (negatively) on the rational status of one’s existing beliefs. To see this, consider conciliationists first. Conciliationists usually support their position by arguing that peer disagreement is higher-order evidence that one’s existing belief is irrational. For example, Christensen says that the peer’s disagreement gives one evidence that one has made a mistake in interpreting the original evidence, and that such evidence should diminish one’s confidence in P ” (2009: 757). Similarly, Feldman concludes that the key fact about disagreement is that the believer’s first-order evidence may support one proposition, while “the disagreement provides evidence that the first-order evidence does not support that proposition” (2009: 305).

Conciliationists are not alone in taking disagreement as higher-order evidence bearing on the rational status of one’s beliefs. Most proponents of more steadfast views also acknowledge that disagreement can constitute higher-order evidence bearing on the rational status of one’s beliefs. For example, Kelly (2010), who is among the first to present disagreement as a kind of higher-order evidence, admits that the higher-order evidence provided by disagreement can have an impact on what one ought to believe. He holds that the mistake many conciliationists make is to say that this impact is always very significant (calling, for example, for a significant reduction in confidence, or suspension of judgment) regardless of what one’s first-order evidence supports. On his contrary “Total Evidence” view of disagreement, the fact of peer disagreement should be added to one’s total stock of evidence as simply one more piece of evidence – in this case “higher-order psychological evidence” (143) – and “what it is reasonable to believe depends on both the original, first-order evidence as well as on the higher-order evidence that is afforded by the fact that one’s peers believe as they do” (142).

¹⁶ Defenders of conciliationism include Christensen (2007), Elga (2007), Feldman (2007, 2009), among others.

¹⁷ Defenders of the steadfast view include Kelly (2005), Van Inwagen (2010), and Titelbaum (2015), among others.

¹⁸ For intermediate views of various different stripes, see, e.g., Enoch (2010), Lackey (2010a, 2010b), Kelly (2010), and Worsnip (2014).

Moreover, proponents of the more radical steadfast view that disagreement is never epistemically significant in and of itself, such as Kelly in his earlier (2005) work, and Titelbaum (2015), rest their view on arguing against the epistemic significance of higher-order evidence more generally. Titelbaum (2015), for example, argues that the disagreeing agent whose view is in fact well-supported by their first-order evidence should retain their belief, because the facts about what their first-order evidence supports are *a priori* and thus (on his view) incapable of being defeated by higher-order evidence, of which Titelbaum assumes the fact of disagreement is an example. Again, this assumes that *vere* disagreement to be epistemically significant, it would be significant *qua* higher-order evidence.

The structure of this dispute between conciliationist and steadfast views shows that despite their divergence, both parties think of disagreement as a kind of (putative) higher-order evidence. Their views only differ in how they think this piece of evidence interacts with other bits of evidence in deciding what we should believe. On this question, they think they have to rely on some general story about how higher-order evidence interacts with first-order evidence. Roughly, conciliationists like Christensen believe that disagreement as higher-order evidence bears on one's belief by playing a special role when interacting with one's first-order evidence, while proponents of more steadfast views like Kelly do not tend to treat this piece of higher-order evidence differently than other first-order evidence, and proponents of extreme steadfast views told that this higher-order evidence has no bearing at all on what one should believe when it comes to first-order propositions.

It is beyond the purpose of this paper to adjudicate between these positions. Instead, in the next subsections, we will zoom in on this commonly accepted claim that disagreement is best understood solely as higher-order evidence, and examine it in light of the definition of higher-order evidence that we offered in the previous section. We aim to show that while disagreement can be thought of as a kind of higher-order evidence, this does not fully describe the evidential status of disagreement. Instead, we propose that disagreement can be both first-order evidence and higher-order evidence bearing on the same proposition.

2.2 *Disagreement as higher-order evidence*

Why is it standardly thought that disagreement is a kind of higher-order evidence? First, let's reconstruct why peer disagreement about whether p is often taken to be evidence that one's doxastic attitude toward p is irrational. As we see it, the argument goes as follows:

P1. If my peer and I share the same body of evidence regarding p and both respond to our evidence rationally, we will arrive at the same doxastic attitude toward p . (**Uniqueness**)

P2. My peer and I did not arrive at the same doxastic attitude toward p . (**Disagreement**)

C1. Therefore, either we do not share the same body of evidence regarding p , or at least one of us must have responded irrationally to our evidence. (From P1, P2)

P3. I share the same body of evidence regarding p with my peer. (**Peerhood 1**)

P4. My peer is not more likely to be irrational than I am. (**Peerhood 2**)

C2. Therefore, there is at least a 50% chance that I responded to my evidence regarding p irrationally. (from C1, P3, P4)

While the argument is valid, it is worth noting that it requires various assumptions to be sound. First, the argument will only apply in cases in which one's disputant is one's epistemic peer, where this involves both sharing the same evidence (P3) and being equally likely to arrive at a rational belief from that evidence (P4). Second, the argument assumes the controversial Uniqueness Thesis. Stated generally, this thesis states that, among incompatible doxastic attitudes toward a proposition, a single body of evidence can only justify one of these attitudes.¹⁹ Applied to the case of disagreement, this entails P1: if both sides of a disagreement evaluate the same body of evidence rationally, they will come to the same (uniquely justified) doxastic attitude.²⁰ This is a crucial assumption of the argument. If Uniqueness is false, then there exist perfectly rational disagreements among those who share the same evidence. As a result, we will not be able to derive the argument's conclusion from the mere fact of peer disagreement.

Finally, note that the conclusion of the argument is, so far, just that there is at least a 50% chance that I have responded to my evidence regarding p irrationally. We've not yet said anything about whether or how this bears on the rational status of my doxastic attitude toward p . So, even if all the premises of the argument hold true, it does not yet follow that disagreement has epistemic significance *qua* higher-order evidence.

To fix ideas, let's suppose that my doxastic attitude toward p is one of belief. Then, if the argument just considered succeeds, it shows that peer disagreement constitutes evidence for the higher-order proposition that my belief that p may well be irrational—call this higher-order proposition *HOP-P*. Given our definition of higher-order evidence, the question now at hand is whether, in virtue of its constituting evidence for *HOP-P*, peer disagreement thereby affects the rational status of one's belief in the first-order proposition p (to remind ourselves that this proposition is a first-order one, let's henceforth switch to labelling it as '*FOP-P*'). If it does, then peer disagreement will be higher-order evidence with respect to *FOP-P*, given our definition.

Recall from section 1 that there are two distinct ways in which higher-order evidence can bear on the rational status of (doxastic attitudes toward) a first-order proposition. First, it can do so by constituting evidence against the first-order proposition in question (albeit derivatively on its constituting evidence against the corresponding higher-order proposition), as is plausible in Expert Detective. Or second, it can do so by affecting one's justification for one's doxastic attitude toward the first-order proposition in some other way (for example, by undercutting one's existing justification for the belief in question), as is plausible in Wishful Thinking.

Insofar as your disagreement indicates that your belief in *FOP-P* may well be irrational, does it thereby also constitute evidence against *FOP-P*? Interestingly, the answer seems to be negative. This is because in general, the fact that a belief is irrational is not a reliable indicator that it is false. For example, suppose it turns out that you believe differently from your peer (in part) because you irrationally overestimate the support your first-order evidence has on a proposition *FOP-P*. This fact

¹⁹ Defenders of the Uniqueness Thesis include Feldman (2007), White (2005), Matheson (2009), among others.

²⁰ To be precise, P1 expresses a slightly more stringent requirement than the general Uniqueness Thesis, because it requires not only that one body of evidence can at most justify one proposition or doxastic attitude, but also that this is an impersonal fact that holds true for *any agent* with that evidence.

indicates that you are no longer justified in believing *FOP-P* (with the confidence that you do), but it does not itself bear on the likelihood that *FOP-P* is true (or false). In the absence of more specific information about your psychological tendencies, we have no reason to suppose that you are more likely to overestimate the evidence for a proposition that is false than for a proposition that is true (or vice versa). So the mere fact that you have misestimated some evidence for a proposition is neither evidence for or against the truth of that proposition. More generally, reasons to think that one's belief was formed by an *unreliable* belief-forming mechanism, i.e. one that does not reliably lead to true belief, are not necessarily reasons to think that it was formed by an *anti-reliable* belief-forming mechanism, i.e. one that reliably leads to *false* beliefs.

Thus, we conclude that if peer disagreement functions as higher-order evidence that bears on the rational status of one's belief in *FOP-P* – that is, derivatively on its support for *HOP-P* – it does so not by constituting evidence against *FOP-P*, but through some other mechanism such as undercutting defeat. In this sense, we might say the discovery of peer disagreement is like the discovery of one's tendency of wishful thinking—as evidence for the relevant higher-order proposition, it can bear on the rational status of one's belief without itself constituting clear evidence for or against the belief's content. Thus, peer disagreement does fit our definition of higher-order evidence (HOE₃) with respect to our first-order doxastic attitudes.

2.3 *Disagreement as first-order evidence?*

However, the foregoing considerations also raise a question: can it really be that peer disagreement provides no evidence against the truth of our beliefs? We've just shown that if peer disagreement only functions as *higher-order* evidence with respect to our beliefs in first-order propositions such as *FOP-P* – that is, derivatively on its status as evidence for higher-order propositions such as *HOP-P* – then it does not bear on the rational status of our beliefs by constituting evidence regarding their truth. But we find this result very counterintuitive. Intuitively, the fact that a generally reliable agent believes *p* is at least some evidence for *p*, and by the same token, the fact that a generally reliable agent believes the negation of a proposition is at least some evidence against that proposition. But (if one is generally reliable), when one learns that one's peer disagrees with one's belief that *p*, one learns that a generally reliable agent believes not-*p*. Thus, it seems like learning that one's peer disagrees with one's belief *should* provide some evidence against one's belief.

This leads us to question the assumption that peer disagreement only provides the kind of higher-order evidence we've been discussing. Could peer disagreement also have some separate rational bearing on our beliefs that is not derivative on its status as evidence for the higher-order proposition discussed above? That is, could it play a separate rational role in addition to its role as higher-order evidence?

Our answer is yes. Our suggestion is that in addition to its role as higher-order evidence, disagreement also plays a straightforward role as first-order evidence.²¹ In the latter role, it functions much like ordinary testimony. Compare a case of disagreement – where one believes *p* and subsequently

²¹ Some philosophers have noted that disagreement could function as first-order evidence, but this is typically only mentioned in passing. See, e.g., Feldman (2009: 305), Christensen (2019: 17), and Titelbaum (2019, sec. 11.6).

learns that one's peer believes not- p – with a case of ordinary testimony – where one has no prior belief about whether p , and subsequently receives testimony from a (person that you justifiably believe to be a) reliable testifier that not- p . In both cases we receive some information about what a reliable epistemic agent believes. It is overwhelmingly plausible that in the case of ordinary testimony, learning that a reliable agent believes not- p is some (regular, first-order) evidence against p .²² But if this is so, it is very hard to see why it wouldn't also be (regular, first-order) evidence against p in the case of disagreement. Admittedly, the case of disagreement is one where one presumably has one's reasons for believing p , reasons that if weighty, may weigh against one's peer's testimony to the contrary. But that doesn't mean that this testimony should suddenly have no (first-order) evidential weight.

Indeed, peer disagreement can be seen as a *special case* of testimony, with two special features: i) the speaker is one's epistemic peer, who shares the same evidence with us, and ii) as the hearer, we already have an existing belief that is incompatible with speaker's belief. We've shown how these special features may give peer disagreement a special evidential status as higher-order evidence that ordinary testimony lacks. Yet for all this, it *also* shares the generic features that ordinary cases of testimony have, and thereby also constitutes first-order testimonial evidence. The fact that peer disagreement has rational bearing on our beliefs as higher-order evidence shouldn't obscure the fact that it also has an independent bearing as testimonial evidence.

This teaches us an important lesson: peer disagreement can have more than one evidential role. One such role is that when you believe p and then you learn that your peer believes not- p , the latter fact is some direct (*pro tanto*) evidence for not- p . Here peer disagreement is functioning as first-order evidence in the usual testimonial way. It bears on one's belief in p by directly constituting evidence against p . The second is that when you believe p and then you learn that your peer (who shares your evidence) believes not- p , that is (at least given the assumption of Uniqueness) some evidence that you (may) have misevaluated the evidence or otherwise made an irrational mistake. Here, given further assumptions about the bearing of such discoveries on the rational status of our first-order beliefs, peer disagreement is functioning as higher-order evidence. It bears on one's belief in p by constituting evidence for a higher-order proposition concerning the (ir)rationality of this belief. As our definition of higher-order evidence makes clear, it is entirely possible for one piece of evidence to be both higher-order and first-order evidence with respect to the very same proposition p . Peer disagreement turns out to be precisely one such case.

3. Upshots for the peer disagreement debate

In this final section, we discuss how the foregoing sheds new light on the debate about the epistemic significance of peer disagreement. Since peer disagreement has two different potential evidential roles,

²² Some philosophers (e.g. Feldman 2007, Eder & Brössel 2019) seem to treat ordinary testimony as higher-order evidence, as they claim testimony is evidence for “the existence of some (specified or unspecified) unknown evidence in the evidential state of another agent that supports p ” (Eder & Brössel 2019: 81) But we disagree with this characterization of testimony. Instead, we treat a reliable person's testifying that p as first-order evidence for p , because we think the fact that any reliable source of information – say, a reliable instrument – says that p is some first-order evidence for p , and *people* are no exception. While it's true that we can infer the existence of some body of evidence from the fact that someone testifies that p , we don't think this adds further evidential support for p on top of the first-order support provided by the testimony.

the explanation of its epistemic significance might go via a story about its role *qua* first-order evidence, or via a story about its role *qua* higher-order evidence, or both. This clarification makes it easier to establish the conclusion that peer disagreement generates conciliatory pressure for one to revise one's beliefs. This is because the explanation of conciliatory pressure that goes via disagreement's role as first-order evidence applies both given fewer controversial theoretical assumptions and in a wider range of cases, compared with that which goes via its role as higher-order evidence.

First, as we saw earlier, the explanation that goes via disagreement's role as higher-order evidence relies on the controversial Uniqueness Thesis, the thesis that one body of evidence can only justify one unique doxastic attitude. Without this assumption, peer disagreement will no longer indicate one's possible irrationality as shown before. But Uniqueness is a controversial thesis that faces many challenges, and some question it precisely because they believe that different agents might disagree with each other without either's rationality being compromised.²³

Second, another assumption that we need to make in order to take the higher-order evidence route is that higher-order evidence does indeed bear on the rationality of one's doxastic attitudes toward first-order propositions. While this is perhaps a less controversial view than Uniqueness, there are certainly philosophers who reject it, and who claim that the evidence that your belief is irrational does not make this belief irrational to hold if the arguments and (first-order) evidence on which one's belief is based are, in fact, objectively good ones.²⁴ But our discussion shows that such a stance, even if it is right, will not suffice to show that disagreement is not epistemically significant. Peer disagreement can still function as first-order evidence against one's beliefs.

Third, as we saw earlier, the explanation of conciliatory pressure that goes via disagreement's role as higher-order evidence applies only in a somewhat limited range of cases, namely that those where one's disputant shares one's body of evidence, and is equally reliable in evaluating it. Indeed, some worry that outside of philosophers' toy examples, cases of genuine epistemic peerhood may be quite rare.²⁵ By contrast, the explanation of conciliatory pressure that goes via disagreement's role as first-order evidence does not require any assumption of peerhood: it applies straightforwardly so long as one's disputant is (justifiably thought to be) a generally reliable person, and regardless of one's own comparative reliability. Consequently, the first-order evidence route can generate conciliatory pressure in most cases of real-world disagreement, insofar as the disputants are generally more reliable than chance.

One way to respond to the rareness of strict peerhood is by loosening the standards for one to rationally identify others as one's epistemic peer. For example, some defenders of the epistemic significance of disagreement acknowledge that actual peerhood might be rare, and that we might not usually have positive evidence for believing that someone shares the exact same evidence with us or is equally reliable as us. But they argue instead that you should revise your beliefs as long as you lack appropriate or clear grounds to think of yourself as epistemically superior than your disputant.²⁶

²³ Cf. Christensen (2016), Schoenfield (2014), and Weintraub (2013), among others. See Ballantyne & Coffman (2012) for an argument as to why peer disagreement might be higher-order evidence despite the failure of Uniqueness.

²⁴ See footnote 9.

²⁵ Cf. Gelfert (2011), King (2012), Vavova (2014).

²⁶ Cf. King (2012: 267-9); Sherman (2015: 431-4); Machuca (2022: 66-8).

However, such strategies face many challenges. For example, some claim that disagreement itself is at least some reason to disregard our disputants, or insist that we do need clear positive evidence of peerhood from our disputant in order to take disagreement with them seriously.²⁷ Perhaps more importantly, when peerhood is understood in this weaker way, peer disagreement provides rather weaker higher-order evidence. After all, the fact that someone who you know relatively little about holds an incompatible belief with you doesn't seem to strongly indicate that you have a large chance of making rational mistakes in forming your belief.

Thus, we think it is better for those who want to generalize the epistemic significance beyond cases of precise peerhood to appeal to the role of disagreement as *first-order* evidence. The first-order route provides a better theoretical framework under which we can establish the conciliatory pressure of disagreement in cases of non-peerhood. The fact that someone with unclear epistemic credentials disagrees with you might not be evidence for your irrationality, but it is certainly some testimonial evidence that your belief is false. Moreover, it remains such evidence even if you are not being irrational, if it is possible that your disputant has access to evidence you lack. Thus, attention to the role of disagreement as first-order evidence has the capacity to shore up the case for broadly conciliatory views about peer disagreement.

References

- Avnir, Y. & Scott-Kakures, D. (2015). "How Irrelevant Influences Bias Belief," *Philosophical Perspectives*, 29: 7-39.
- Ballantyne, N. & Coffman, E.J. (2012). "Conciliationism and Uniqueness," *Australasian Journal of Philosophy*, 90/4: 657–670.
- Begby, E. (2021). "The Problem of Peer Demotion, Revisited and Resolved," *Analytic Philosophy*, 62/2: 125-140.
- Christensen, D. (2007). "Epistemology of Disagreement: The Good News," *The Philosophical Review*, 116(2), 187-217.
- (2009). "Disagreement as Evidence: The Epistemology of Controversy," *Philosophy Compass*, 4/5: 756-767.
- (2010). "Higher-Order Evidence," *Philosophy and Phenomenological Research*, 81/1: 185-215.
- (2016). "Conciliation, Uniqueness, and Rational Toxicity," *Noûs*, 50/3: 584–603.
- (2019). "Formulating Independence," in Skipper & Steglich-Petersen (eds.), *Higher-Order Evidence: New Essays*. Oxford University Press.
- Comesaña, J. (2012), "Conciliation and Peer-Demotion in the Epistemology of Disagreement," *American Philosophical Quarterly*, 49: 237–251.
- Eder, Anna-Maria A. & Brössel, P. (2019). "Evidence of Evidence as Higher Order Evidence," in Skipper & Steglich-Petersen (eds.), *Higher-Order Evidence: New Essays*. Oxford University Press.
- Elga, A. (2007). "Reflection and Disagreement," *Noûs*, 41/3: 478–502.
- Enoch, D. (2010). "Not Just a Truthometer: Taking Oneself Seriously (but not Too Seriously) in Cases of Peer Disagreement," *Mind*, 119/476: 953-997.

²⁷ Cf. Comesaña (2012), Vavova (2014), Begby (2021).

- Feldman, R. (2005). "Respecting the Evidence," *Philosophical Perspectives*, 19: 95-119.
- (2007). "Reasonable Religious Disagreements," in Antony (eds.), *Philosophers without Gods*. Oxford University Press.
- (2009). "Evidentialism, Higher-Order Evidence, and Disagreement," *Episteme*, 6/3, 294-312.
- Gelfert, A. (2011). "Who is an epistemic peer?," *Logos and Episteme*, 2/4: 507–514.
- Hedden, B. & Dorst, K. (forthcoming). "(Almost) All Evidence is Higher-Order Evidence," *Analysis*.
- Horowitz, S. (2014). "Epistemic Akrasia," *Noûs*, 48/4: 718-744.
- González de Prado, J. (2020). "Dispossessing Defeat," *Philosophy and Phenomenological Research*, 101/2: 323-340.
- Kelly, T. (2005). "The Epistemic Significance of Disagreement," *Oxford Studies in Epistemology*, 1: 167–196.
- (2010). "Peer Disagreement and Higher-Order Evidence," in Feldman & Warfield (eds.), *Disagreement*, Oxford University Press.
- King, N. (2012). "Disagreement: What's the problem? Or a good peer is hard to find," *Philosophy and Phenomenological Research*, 85/2: 249–272.
- Lackey, J. (2010a). "A Justificationist View of Disagreement's Epistemic Significance," in Haddock, Millar & Pritchard (eds.), *Social Epistemology*. Oxford University Press.
- Lackey, J. (2010b). "What Should We Do When We Disagree?," *Oxford Studies in Epistemology*, 3: 274-293.
- Lasonen-Aarnio, M. (2014). "Higher-Order Evidence and the Limits of Defeat," *Philosophy and Phenomenological Research*, 88/2: 314-345.
- Lord, E. & Sylvan, K. (2021). "Suspension, Higher-Order Evidence, and Defeat," in Brown & Simion (eds.), *Reasons, Justification, and Defeat*. Oxford University Press.
- Machuca, D. (2022). "Personal Information as Symmetry Breaker in Disagreements," *Philosophy*, 97/1: 51-70.
- Matheson, J. (2009). "Conciliatory Views of Disagreement and Higher-Order Evidence," *Episteme*, 6/3: 269–279.
- Pollock, J. (1986). *Contemporary Theories of Knowledge*. Rowman and Littlefield.
- Schoenfield, M. (2014). "Permission to Believe: Why Permissivism is True and What It Tells Us About Irrelevant Influences on Belief," *Noûs*, 47/1: 193–218.
- Sherman, B. (2015). "Unconfirmed Peers and Spinelessness," *Canadian Journal of Philosophy*, 45/4: 425–444.
- Skipper, M. & Steglich-Petersen, A. (2019). "Introduction," in their *Higher-Order Evidence: New Essays*. Oxford University Press.
- Staffel, J. (2021). "Pro Tem Rationality," *Philosophical Perspectives*, 35/1: 383-403.
- Titelbaum, M.G. (2015). "Rationality's Fixed Point (or: In Defense of Right Reason)," *Oxford Studies in Epistemology*, 5: 253-294.
- (2019). "Return to Reason," in Skipper & Steglich-Petersen (eds.), *Higher-Order Evidence: New Essays*. Oxford University Press.
- van Inwagen, P. (2010). "We're Right, They're Wrong," In Feldman & Warfield (eds.), *Disagreement*, Oxford University Press.

- Vavova, K. (2014). "Moral Disagreement and Moral Skepticism," *Philosophical Perspectives*, 28: 302-333.
- Weatherson, B. (2019). *Normative Externalism*. Oxford University Press.
- Weintraub, R. (2013). "Can Steadfast Peer Disagreement Be Rational?," *Philosophical Quarterly*, 63/253: 740–759.
- White, R. (2005). "Epistemic Permissiveness," *Philosophical Perspectives*, 19: 445–459.
- (2010). "You Just Believe That Because...", *Philosophical Perspectives*, 24: 573-615.
- Worsnip, A. (2014). "Disagreement about Disagreement? What Disagreement about Disagreement?," *Philosophers' Imprint*, 14/18.
- (2018). "The Conflict of Evidence and Coherence," *Philosophy and Phenomenological Research*, 96/1: 3-44.