

More Hope for Conciliationism

Martin Justin

University of Maribor, Faculty of Arts, Maribor, Slovenia

0000-0002-2614-9423

martin.justin1@um.si

Accepted for publication in *Episteme*

Abstract

The view that epistemic peers should conciliate in cases of disagreement—the Conciliatory View—had been an important view in the early days of the peer disagreement debate. Over the years, however, the view has been the target of severe criticism; an “obituary” was already written for the view, and, as a recent proclamation has it, there is “no hope” for it. In this paper, I will argue that we should keep the hope alive by defending the Conciliatory View of peer disagreement. The primary strategy of my defense will be to separate the claims made by the view specific to peer disagreement and claims that concern higher-order evidence more generally. This separation allows us to see which problems cannot be addressed in the context of peer disagreement alone. As I will argue, the upshot of making this distinction is that although the jury is still out on whether higher-order evidence should affect our first-order doxastic states, the Conciliatory View likely follows if it does.

Keywords: peer disagreement, higher-order evidence, the conciliatory view, epistemology of disagreement, self-undermining.

1 Introduction

One of the tasks taken up recently by epistemologists is to formulate a modest account of epistemic rationality: an account that would reckon with the fact that we are fallible thinkers (see, for example, Christensen, 2020; Dorst, 2020). In large part, this task consists of explaining what role (if any) we should assign to evidence that our doxastic states might be flawed. Such evidence, usually called “higher-order evidence”, comes in different forms—from learning about a common cognitive bias to discovering that you were given a reasoning-distorting drug (for reviews of higher-order evidence, see Horowitz, 2022; Ye, 2022)—but it poses a similar problem in most cases. On the one hand, it seems intuitive that getting such evidence should prompt us to revise our doxastic states significantly. On the other hand, however, we still have first-order evidence which strongly supports our original position. Whichever way we go—adopt a new state or stick to the current one—we seem to be ignoring some piece of our total evidence.

One of the most frequently discussed examples of this puzzle is the problem of peer disagreement (for a review, see Frances and Matheson, 2018). It concerns a specific kind of higher-order evidence: learning that someone in a similar epistemic position disagrees with you. Early in the peer disagreement debate, different philosophers have proposed the so-called Conciliatory View which states that getting evidence about such disagreement about some proposition p should prompt you to significantly lower your confidence in your belief that p , perhaps even to the point of suspending judgment (Bogardus, 2009; Christensen, 2007; Elga, 2007; Feldman, 2005, 2006; Kornblith, 2010; Matheson, 2009).

In the years since, the Conciliatory View received largely negative attention. Opponents of the view have argued, among other things, that it relies on false epistemic principles (Enoch, 2010; Lackey, 2010; Lord, 2014), violates true ones (Titelbaum, 2015), has undesirable consequences such as unwanted skepticism (Kelly, 2005), requires reasoners to ignore evidence (Kelly, 2010; Sliwa and Horowitz, 2015), gives wrong, incomplete or pernicious advice (Cruz and Smedt, 2013; Douven, 2010; Jehle and Fitelson, 2009; Rasmussen et al., 2018) or is self-undermining (Christensen, 2020; Dixon, 2024; Elga, 2010; Matheson, 2015). Some have been

quite stark in their condemnation: Lord (2014) has written an “obituary” for the view, while Dixon (2024) has recently stated that there is “no hope” for it.

In this paper, I will argue that we should keep the hope for the Conciliatory View alive by providing a novel defense of it. I will provide a precise account of the Conciliatory View and show that it has the resources to answer the existing attacks. That is not to say that the view presented here is a fully fledged and final account of peer disagreement that can withstand all possible objections. Rather, my more modest goal is to show that the view presented here can be defended against the more pressing problems and can thus be considered as a very much alive working hypothesis about peer disagreement.

My primary strategy will be to separate the claims made by the view specific to peer disagreement and claims that concern higher-order evidence more generally. The relation between the higher-order evidence and peer disagreement debates has been frequently observed in the literature Christensen (2010), Feldman (2009), and Matheson (2024). However, this paper is, to my knowledge, the first to systematically discuss how specific claims that make up the Conciliatory View relate to the broader question of higher-order evidence.

This approach has two significant consequences. First, it allows us to understand better in what ways peer disagreement is independently interesting from the questions about higher-order evidence. Relatedly, it helps us see which problems cannot be addressed in the context of peer disagreement alone. As I will argue, the upshot of making this distinction is that although the jury is still out on whether higher-order evidence should affect our first-order doxastic states, the Conciliatory View likely follows if it does.

The paper is organized as follows. I will present and motivate the Conciliatory View in the next section. Sections 3, 4, and 5 will, in turn, defend specific claims that make up the view. These sections will provide novel defenses against some of the attacks on the view that have not been sufficiently addressed before. Section 6 will defend the general approach taken up in this paper.

2 The Conciliatory View of Peer Disagreement

Consider the following case:

Copy-Editing: You and Rory work as copy-editors at a magazine. You are both considered to be very good at your job, and as far as you know, you have a similar track record. You know this because the magazine you work for follows a rigorous styling guide, so there is usually only one correct way to copy-edit a given text. One day, Rory and you are assigned to work on the same article. After you finish, you turn to Rory, who works at the desk beside you, and proclaim: “There were 76 mistakes in the text; I’m quite confident.” Rory looks at you and replies: “That can’t be right! I only counted 70.” You both recount the corrections and confirm that there is indeed a disagreement: you think the writer made 76 mistakes, while Rory only found 70. You are well rested and were focused while working, and as far as you can tell, the same goes for Rory.¹

How should one respond to such disagreements? Should you change your confidence in your belief, or should you stick to it? Intuitively, the former seems to be the case, at least in the case above. Given the strict styling guide, either you or Rory (or both) must be mistaken about the number of mistakes in the text. However, both of you are equally good at your job, you have a similar track record at these things, and neither of you is drunk, distracted, or in any other way cognitively impaired at the moment. In other words, neither of you is in an obviously better (or worse) position to get the number of mistakes right. Therefore, you don’t have a reason to prefer your belief over Rory’s. Learning about such disagreement should thus prompt you to *conciliate*—that is, to significantly lower your confidence in the belief that there are 76 mistakes in the text, perhaps even prompting you to suspend judgment about the question.

¹This case is modeled on the canonical ‘restaurant case’ first proposed by Christensen (2007, p. 193). The idea of using copy-editing as an example of disagreement was suggested to me by Dunja Šešelja in a discussion.

Cases like these (and the reasoning behind them) motivate the Conciliatory View of peer disagreement. The view can now be further specified as a conjunction of the following four claims:

Accuracy-Peerhood: Two agents are epistemic (accuracy) peers about p if they are in an equally good epistemic position about the accuracy of their beliefs about p .

Independence: Reasons to discount evidence of your (un)reliability about p (in this case, the fact of peer disagreement) must be independent of your reasoning from first-order evidence to p and your belief about p that you adopted based on this reasoning.

Equal Weight: When involved in an actual disagreement with a peer, both your and your peer's beliefs about p should be given equal weight.

First-Order Import: Evidence of your (un)reliability about p (in this case, the fact of peer disagreement) has a first-order bearing on your belief about p .

This explication of the view closely follows Frances and Mathenson (2018) and Matheson (2024), who call it the Equal Weight View. However, there are some crucial differences. In the remainder of this section, I will elaborate on these differences and provide additional details about the claims.

First, Accuracy-Peerhood can be further specified in two different ways. According to Elga (2007), two agents are epistemic peers if they are equally likely to be correct about the thing in question. Alternatively, Kelly (2005) or, more recently, Frances and Mathenson (2018) define peerhood as being equal in several epistemically relevant factors. Kelly mentions access to relevant evidence, intelligence, thoughtfulness, and freedom from bias. In addition, Frances and Matheson add relevant background knowledge, time devoted to answering the question, intellectual virtues, and freedom from distractions. Some authors also talk about peers being evidential (Christensen, 2007) or cognitive equals (Lackey, 2010) or both (Cruz and Smedt, 2013), which I understand as saying that they are equal in those epistemic factors—evidence, cognitive capacities, or both.

Another distinction is between accuracy- and rationality-peerhood (Christensen, 2016). Accuracy-Peerhood, the claim above, is limited to the accuracy of agents' beliefs. Alternatively, it could concern their rationality—you and Rory, in the above example, might be in an equally good position to both have *rational* beliefs about the number of mistakes in the text. However, this rationality account of peerhood would not get the Conciliatory View off the ground. The view is motivated by the idea that evidence of disagreement indicates a mistake has been made. Since there is only one way to correctly count the errors in the text (as per assumption), it is not possible that both you and Rory are right. In contrast, it's entirely feasible that both of you are rational. Perhaps you both subscribe to subjective Bayesianism and think that the only two constraints on rationality are conditionalization and laws of probability. If we assume that you both obey these two constraints and that the differences in beliefs come from your different priors, both beliefs about the number of mistakes in the text would be rational. Consequently, the Conciliatory View of disagreement must be limited to the cases of disagreement between accuracy peers.

In other words, this way of understanding peerhood evades subscribing to an implausibly strong version of what has been called Uniqueness. Uniqueness, in its strongest version, states that there is at most one maximally rational response to any given batch of total evidence (Christensen, 2016). As Christensen points out, there are good reasons to think this claim is too strong. For one, it is rejected by the aforementioned subjective Bayesianism, a prominent account of rationality. We can weaken Uniqueness by limiting it to accuracy: there is at most one maximally accurate credence in p to any batch of total evidence about p . This claim is much more plausible. In the Copy-Editing case, it is true by assumption, but it is natural to think that it holds in various epistemic domains, especially when we make descriptive claims about the world. The upshot here is that Accuracy-Peerhood only assumes this weaker, more plausible reading of Uniqueness.

Second, Independence prohibits specific reactions to peer disagreement (and other similar evidence) that seem inappropriate at face value since they lead to “blatantly question-begging dismissals of the evidence” (Christensen, 2011, p. 2). For example, in the Copy-Editing case, it prohibits you from saying something like: “But according to my corrections, there are 76 mistakes in the text! Rory,

you're mistaken." Independence is further refined and generalized in contrast to its alternative presented by Frances and Mathenson (2018). On the one hand, it specifies what should be put aside when considering whether to disregard the opinion of a disagreeing peer. On the other hand, it generalizes the claim to other kinds of evidence about our (un)reliability in assessing (or collecting) evidence (for more about the role of Independence in accounts of higher-order evidence, see Christensen, 2019; Ye, 2022).

The distinction between Equal Weight and First-Order Import is perhaps the most significant departure from the existing presentations of the Conciliatory View. The distinction stems from the idea that the introduction of higher-order evidence requires a "two-tiered" picture of doxastic states that distinguishes between our conviction (or justification) that p and our conviction (or justification) that our belief that p (different kinds of this distinction are introduced by Dorst, 2020; Henderson, 2022; Lasonen-Aarnio, 2014; Ye, 2022). Equal Weight claim concerns our higher-order conviction: we should be equally convinced that our belief about p and our peer's belief about p . On the other hand, First-Order Import concerns the relation between the levels or tiers: our first-order conviction should follow our higher-order conviction.

Equal Weight can be motivated by an appeal to the symmetry between epistemic peers. However, it can also be motivated by pointing out the problems of not giving equal weight in cases of peer disagreement. Elga (2007), for example, argues that assigning extra weight to your beliefs in cases of peer disagreement can lead to holding irrational, even absurd, convictions. Imagine that you and Rory independently copy-edit the same long series of articles. You do not have any outside information about the number of mistakes in the texts, but you can compare each other's judgments. Suppose that every time you disagree, you should be 70% confident that you are correct. At the end of the series, you should thus end up extremely confident that you have a better track record at copy-editing than Rory. Elga (2007, p. 487) rightly points out that this result is absurd.

On the other hand, First-Order Import can be motivated by the intuition that our beliefs should at least broadly follow our convictions about what is rational for us to believe in a given situation. More specifically, rejecting First-Order Import

can lead to epistemic akrasia. In this attitude, you “think you ought not to have a certain doxastic attitude towards a proposition p , but you have that attitude towards p anyway” (Henderson, 2022, p. 518).

The following three sections of the paper will defend Independence, Equal Weight, and First-Order Import, respectively. Although some authors contest that epistemic peers can be found in real life (King, 2012), I will agree with Cocchiaro and Frances (2021) that Accuracy-Peerhood can be applied to cases of actual disagreement and leave further discussion of this issue aside. There are two reasons for this. First, the other aspects of the Conciliatory View are much more contentious and thus need more attention. Second, as Matheson (2024) points out, cases of peer disagreement are interesting even if they are complete fictions. He writes that idealized cases of peer disagreement that focus on epistemically relevant features provide us with a starting ground for approaching messier cases of real-world disagreements.

3 Defending Independence

This section will defend Independence. First, I will discuss objections by Kelly (2013), Lackey (2010), and Lord (2014). These objections argue that Independence sometimes gives intuitively wrong answers and that no general appropriately weakened principle can be constructed. I will present a novel reading of independence that avoids these counter-examples. Second, I will discuss some more recent counter-examples due to Moon (2018) and show that my reading of Independence can deal with them without Moon’s suggested amendments.

Let’s first look at the argument presented by Kelly (2013). Consider the following case:

The Black Death: After finishing copy-editing the text, Rory goes to the cafeteria for lunch. There, she meets her colleague, editor Max. They make small talk, and Rory mentions that she recently read a fascinating article about the 14th century plague pandemic in the New Yorker. Max listens to her and then announces: “I don’t believe that that

pandemic ever happened.” Normally, Rory would consider Max a great editor and her epistemic peer regarding common historical knowledge; in addition, Max does not seem especially tired, drunk, or cognitively impaired in any other way.²

Should Rory revise her belief that the black death happened after she learns that Max, her epistemic peer, disagrees with her? According to the Conciliatory View, she should. But that seems highly unintuitive: she has no other reason to doubt the existence of the plague pandemic, so surely she should rather discount Max’s opinion than change hers.

As Kelly points out, this strategy seems to be prohibited by Independence. According to the claim, Rory’s reasons for discrediting Max should be independent of her reasoning about the thing in question. However, in cases such as The Black Death, it is precisely Rory’s overwhelming evidence for the existence of the pandemic that most naturally serves to discredit Max. In other words, it doesn’t seem unreasonable for her to say something like: “Max says the Black Death didn’t happen. But it *did*. Thus, Max must be wrong about this.” Thus, The Black Death is a counter-example to Independence.

One way to defend Independence, employed by Christensen (2011), is to show that Rory does have other independent reasons to discredit Max. In cases like Copy-Editing, the most likely explanation for the disagreement is that one of the peers has made a mistake. Given that they are in an equal epistemic position concerning the proposition, the peers have no reason to privilege one or the other’s beliefs. But in cases of more extreme or unusual disagreements, it becomes exceedingly unlikely that a mistake has been made. It seems hard to imagine that Max or Rory are mistaken about the bearing of the available evidence on the proposition that the medieval plague pandemic happened. There are other much more likely explanations: one of them is joking or being dishonest about her beliefs, one of them is under the influence of drugs, one of them has some ulterior motives (for example, has lost a silly bet and now has to pretend), etc. Thus, Rory can use all of these explanations to discredit Max’s opinion, and none has anything to do with her

²This case is based on the Holocaust Denier case, provided by Kelly (2013, p. 40).

belief about the black death.

Kelly (2013) considers and dismisses this defense. While I do not necessarily agree with Kelly's argument against Christensen's defense,³ I do want to suggest an alternative defense of Independence. This defense recognizes that Independence is targeted—it does not necessarily work as a blanket restriction on all reasoning from evidence to the proposition in question. Specifically, it forces us to bracket only reasons directly put into question by the higher-order evidence we received.

We can make this explicit by rephrasing the Independence principle as follows:

Independence*: Reasons to discount higher-order evidence of your (un)reliability about p (in this case, the fact of peer disagreement) must be independent of your reasoning from first-order evidence to p *that your higher-order evidence brackets* and from your belief about p that you adopted based on this reasoning.

Let's take a closer look at how we should understand this added restricting clause. Per assumption, Rory believes Max is her peer concerning common historical knowledge. In other words, if p is a proposition about a commonly known historical event, Max and Rory are in an equally good position to have accurate beliefs about p . What does this tell us about the scope of Independence*, i.e., the limits it poses on Rory's reasons to dismiss disagreements about commonly known historical events with Max? One option is to say that it prevents her from using all possible reasoning that leads from her first-order evidence to beliefs about commonly known historical

³Kelly argues that although possible, the reasoning described by the strategy is too artificial to be attributed to real reasoners. The idea that we need this whole procedure of finding independent reasons to discredit someone seems to him "akin to the suggestion that, in a case in which I discover that I hold two inconsistent beliefs, I should evaluate the credentials of one belief while bracketing my assessment that I have overwhelming evidence for the other." Although I agree with Kelly that the reasoning described by Christensen (2011) is not the most natural, I don't find this a good counter-argument. Reasoning that violates Independence might sometimes feel more natural, but it is nevertheless suspiciously question-begging. Therefore, we should prefer alternatives. In addition, his comparison does not seem apt: peer disagreement is more similar to arriving at two inconsistent conclusions about the same proposition based on comparable evidence at two different times than coming to recognize (at one time) that two of my beliefs are inconsistent and supported by different evidence. Unlike Kelly's case, in these cases, it seems unreasonable to discard the second conclusion because it conflicts with the first. I thank the anonymous reviewer of this journal for helping me clarify this point.

events. This is a strong reading, which gives an unintuitive answer in The Black Death case. Since the Black Death is a commonly known historical event, Rory is prohibited from using anything that pertains to or could be used to support her belief that the plague pandemic happened as a reason to discount the disagreement as insignificant. Since, per assumption, she also lacks independent reasons (e.g., Max is not visibly drunk), this reading of Independence* implies Rory should conciliate.

Fortunately, a weaker reading of Independence* is available to us. This one says that if two agents are accuracy peers with regard to commonly known historical events, the reasons to dismiss their disagreement must be independent of reasoning from this common historical knowledge. If p is a proposition about a commonly known historical event (e.g., “the Black Death happened”) and Rory and Max are peers with regard to commonly known historical events, Independence* now prohibits them to refer to commonly known historical evidence that supports p , in discounting their disagreement. In the above case, it prohibits Rory from arguing: (1) The Black Death is a commonly known historical event, I learned about it in junior high! (2) Max believes the Black Death didn’t happen. (3) Thus, Max is wrong about this event.

However, in contrast to the above stronger reading, it does not prohibit other lines of reasoning that still refer to beliefs about commonly known historical events but use other evidence. For example, (1) I got reliable testimony that the Black Death happened, which goes beyond commonly shared historical knowledge. (2) Max believes the Black Death didn’t happen. (3) I’m not sure Max has any knowledge about the Black Death beyond what’s commonly known. (4) Thus, Max is wrong about the Black Death. Since Rory got a testimony that the medieval plague pandemic happened (she read the article), she could use this reasoning to dismiss her disagreement with Max. On the other hand, this weak reading still gives the correct responses in cases where we want the peers to conciliate. For example, in the Copy Editing case, you are prohibited from using any copy-editing-related reasons to discredit Rory.

One might worry that this reading of Independence* is too weak. It only says that reasons we provide for dismissing disagreements (or other kinds of higher-order evidence) must be independent of reasoning that explicitly refers to subjects covered

by peerhood (or higher-order evidence). For example, Rory could argue: (1) I *feel* like the medieval plague pandemic happened. (2) Max believes that it didn't. (3) Thus, Max is wrong about the pandemic. This line of reasoning, indeed, is not prohibited by the weak reading of Independence* presented above. But that does not mean it gives Rory a good reason to disregard Max. There are other grounds to dismiss it. We do not usually take feelings as sufficient justification for our beliefs. It is not clear why we should in this case. Therefore, although it does not violate Independence*, the above argument is not a valid reason to dismiss the disagreement. It is not the role of Independence* to filter all invalid reasons for dismissing higher-order evidence. Its role is to block specific reasons—suspiciously question-begging reasons—while other epistemic principles can block other non-question-begging but bad reasons.

To sum up the preceding discussion, I argued that Independence is targeted, forcing us to bracket only reasons directly put into question by the higher-order evidence we received. There are two ways to read this. According to a stronger reading, this limited principle, which I called Independence*, blocks every reason that pertains to the peer proposition in question. In other words, if Max and Rory are equally likely to have accurate beliefs about commonly known historical events and p is such event, their reasons to discredit each other should be independent of anything that pertains to p . I rejected this in favor of a weaker reading. This one states that Independence* blocks only some reasons pertaining to the peer proposition. Specifically, if Max and Rory are equally likely to have accurate beliefs about commonly known historical events and p is such event, their reasons to discredit each other should be independent only from other commonly known historical facts that support p . In contrast, if Max and Rory were peers but also renowned experts on the Black Death, Independence* would have a much broader scope since their expertise includes much more than just common historical knowledge.

This reading of Independence can also evade critiques that were leveled against Christensen's strategy. One such objection to Christensen can be found in Lord (2014). Lord points out that "Christensen's strategy is to show that Independence is compatible with stories that deliver the right results." He concedes that Christensen can provide a story in which Rory would be justified in dismissing Max without

violating Independence. However, Lord finds this insufficient; to him, Independence is false because it blocks other legitimate stories that give the correct result. For example, the one where Rory discredits Max based on her overwhelming evidence of the plague pandemic. This objection applies to Christensen's strategy. Christensen insists that Rory indeed cannot use her evidence but must instead build alternative explanations of the disagreement, e.g., that Max is joking. However, it does not apply to the strategy used here. We already saw that Rory can use some of her evidence to discredit Max. Some details of the account are still missing, of course. For example, it is not entirely clear what evidence Rory possesses should be considered as part of the common historical background and what not. Does something she vividly remembers from school count as independent evidence for the event, or should it be discarded as part of the common historical background? Questions like this remain open.⁴

Nevertheless, I think my account, in general, gives us enough space to answer Lord's criticism of Independence. As stated, my goal is not to provide a fully fledged account of Independence and of the Conciliatory View that will decidedly answer all possible criticism. Rather, I aim to show that the view can, in principle, be defended and there is hope that it will be fully vindicated in the future. Pointing to the underappreciated aspect of Independence as a targeted rather than a blanket restriction does just that—although some questions remain, it diffuses the main

⁴There are two additional worries about the reading of Independence* defended here. The first concerns a potential conflict between Independence* and Accuracy Peerhood. The worry is that given that Rory has an additional piece of evidence Max lacks, they are not accuracy peers anymore. If this is the case, Independence* avoids counterexamples at the cost of sacrificing peerhood. The conflict indeed arises with accounts of peerhood that require strict evidential equality. However, Accuracy Peerhood, as stated above, allows for weaker accounts that do not conflict with Independence*. For example, one such account could state that if Rory and Max are accuracy peers with regard to *p*, they are, on balance, equally likely to have accurate beliefs about *p*, without implying that they are strict evidential or cognitive equals. This reading then leaves open the option that, in a specific case, one of them has some advantage without breaking the peerhood. A fully developed account of the Conciliary View should further specify exactly what reading of Accuracy Peerhood is compatible with Independence*. The second worry is that even the weaker reading of Independence* still gives wrong answers in cases similar to The Black Death. What if, for example, Max also reads the same magazine article about the pandemic? If we assume that Max and Rory are peers with regard to reading comprehension, which is reasonable since they are both educated adults, Rory would be prohibited from dismissing their disagreement by appealing to this piece of evidence. A full account of Independence* should also answer the other questions that remain open. I'm grateful to the reviewer for flagging these further objections to Independence*.

charge of Lord's criticism and allows us to address additional objections.

One such additional objection to Christensen's version of Independence, presented by Lackey (2010), relates to this last point about the ambiguity involved in considering which reasons Independence blocks. Lackey argues that in cases of real-world disagreements, especially if they are unusual or extreme, there is often a considerable amount of ambiguity involved in determining whether someone is a peer. Thus, a balancing act is required. According to Lackey, when considering whether to dismiss a disagreement, one is balancing the support or justification one has for a belief in question with the confidence that the disagreeing person is, in fact, one's peer. Furthermore, Lackey argues that this act of balancing is all there is to determining the epistemic significance of a specific case of peer disagreement. In other words, she accepts this case-to-case balancing wholeheartedly and rejects the idea that we can say something principled, like Independence, about cases of disagreement.

While Lackey's argument is convincing against the view of Independence defended by Christensen (2010) that requires us to consider which explanation of the disagreement is most likely, the view presented here answers it. Specifically, it shows that no consideration of the degree of justification we have for our views is involved in considering whether to dismiss a disagreeing peer. Instead, whether to dismiss a disagreement entirely depends on the scope of higher-order evidence. As already mentioned, some ambiguity can be involved in determining this, especially in cases of peer disagreement. However, the view defended here gives strikingly clear answers in some other cases. Consider this one:

Hypoxia: Rory and her co-pilot Max are out flying their small, unpresurized airplane, wondering whether they have enough fuel to make it to Hawaii. Rory looks at the gauges, dials, and maps and obtains some evidence, *E*, which she knows strongly supports (say to degree .99) either the proposition that they have enough gas (*g*) or that they do not (*not-g*). Thinking it over and performing the necessary calculations, Rory becomes highly confident that *g*; in fact, this is what *E* supports. She announces this to Max. But Max disagrees. Max points out that

given the altitude at which they are flying, they are at moderate risk of hypoxia, a condition that impairs one's reasoning while leaving the reasoner feeling perfectly cogent and clear-headed. Hypoxic pilots performing the kinds of calculations Rory just did only reach the correct conclusion 50% of the time.⁵

Should Rory dismiss Max's worry that they are suffering from hypoxia based on her high confidence that she has enough fuel? Lackey's account calls for this option, but I do not think that is right. The disagreement with Max presents Rory with higher-order evidence that puts into question her reasoning, especially the kind of calculations she needs to perform to determine the amount of fuel left. Assuming Rory has no independent reasons to dismiss Max's worry about the risk of hypoxia, First-Order Import suggests that Rory should revise her belief that *g*. However, a defender of Lackey's view could argue that since Rory is only at moderate risk of hypoxia while she is highly confident of her belief that *g*, she could dismiss the hypoxia worry based on these grounds. Again, I don't think this is the correct analysis here. The reliability of higher-order evidence (i.e., the risk of hypoxia) might have a role in considering the import of higher-order evidence on first-order beliefs. In contrast, taking the justification we have for our beliefs into account here would amount to question-begging reasoning of the form: (1) It is possible that Max is right and I'm unreliable in making statements whether *g*. (2) But I'm quite certain that *g*. (3) Thus, I can remain certain that *g*. Consequently, I think the justification we have for our beliefs should not play a role in considering whether to dismiss evidence about our unreliability.

The final upshot of Independence* is that it gives correct answers in a series of cases, presented by Moon (2018). Consider this one:

Party: Rory is going to a party after work. She learns from a reliable source that David is at a party and that Peter's beliefs about David's location are unreliable. From this conjunction, she infers that David

⁵This is a modified version of a case adapted from Horowitz (2022)—the original case does not include the element of peer disagreement. As far as I can tell, it originates with Elga (2008) and is a standard example in the debate about higher-order evidence. I'm indebted to the anonymous reviewer for suggesting the use of the Hypoxia case here.

is at the party. At the party, she meets Peter, who tells her that David stayed home. From the conjunction she learned before the party, she infers that Peter's beliefs about David's whereabouts are unreliable. Thus, she dismisses Peter's opinion and remains steadfast in her belief that David is at the party.⁶

I think it is natural to see Rory's reasoning in Party as correct. Notice, however, that it seems to violate Independence. In evaluating the epistemic status of Peter's belief about David's whereabouts, Rory relied on the conjunction she learned before the party. However, she used the exact conjunction as a premise in her reasoning behind her belief that David is at the party. In other words, the reasoning she used to discredit Peter was the same as the reasoning behind her belief, with which Peter disagreed.

Based on the Party case, for example, Moon (2018) argues that an independence principle should be qualified with additional conditions. His principle is worded slightly differently, but since he considers it in a more general context, it is also applicable to peer disagreement.

Moon's Independence: "In evaluating the epistemic credentials of another's expressed belief about P, in order to determine how (or whether) to modify my own belief about P, I should do so in a way that neither relies on the reasoning I used to form my initial belief about P nor relies on my belief about P itself, unless

(a) A justifiedly believed premise in S's reasoning about P makes a claim about the epistemic credentials of the other's expressed belief about P [...]."⁷ (Moon, 2018, p. 72)

However, I think this additional condition (a) is not necessary with the reading of Independence* defended here. To see why, consider the scope of Rory's and Peter's peerhood. We can assume that they are peers concerning David's whereabouts.

⁶This case closely follows Counterexample I from Moon (2018, p. 68).

⁷Moon's Independence principle includes two additional such conditions that answer other counterexamples he presents. I will omit both the additional examples and conditions here for the sake of brevity.

Independence* thus prohibits Rory from using David's whereabouts as a premise in her reasoning to discredit Peter. Does this prevent her from using the conjunction "David is at a party, and Peter's beliefs about David's location are unreliable" as such? I think not. The conjunction has a first-order bearing on Rory's beliefs about Peter's reliability, independent of its first-order bearing on her belief about David's location. Since Peter and Rory are peers only regarding the claim that David is at the party, the first-order bearing on Rory's evidence on her beliefs about Peter's reliability is unrestricted by Independence*.

To see this distinction more clearly, consider the following modification of Party:

Party*: Rory is going to a party after work. She learns from a reliable source that David is at a party and that this is a no-drugs party. From this conjunction, she infers that David is at the party. At the party, she meets Peter, who tells her that her drink contains a drug that makes her utterly unreliable in reasoning from conjunctions. From the conjunction she learned before the party, she infers that there are no drugs at the party. Thus, she dismisses Peter's testimony and remains steadfast in her belief that David is at the party.

In this case, it seems much more unreasonable for Rory to remain steadfast in her belief about David's whereabouts. To see why, consider the scope of her higher-order evidence. Learning that she took the drug concerns all her reasoning from conjunctions. In other words, it puts into question all reasons she could possibly derive from the conjunction "David is at a party, and this a no-drugs party.". Thus, she has no independent reasons to dismiss the higher-order evidence of her unreliability. Independence* works as intended.

4 Defending Equal Weight

This part of the paper will take a closer look at Equal Weight. This claim states that when considering how to respond to a disagreement, your and your peer's beliefs should be given equal weight. I will defend this claim against two attacks. One, due

to Enoch (2010), states that the appropriate strategy of responding to peer disagreement can have the effect of giving extra weight to your view. Another, presented by Lasonen-Aarnio (2013), says that Equal Weight relies on an unmotivated and undesirable constraint on our higher-order convictions.

Turning first to Enoch's argument. He first makes a general observation that to get cases of peer disagreement, such as Copy Editing, off the ground, agents involved should have justified beliefs about who their peer is. I think this is right. Consider, for example, that in Copy-Editing, you would form, independently from your disagreement, an unjustified belief that Rory is, in fact, not your peer. Let's say you tossed a coin: since it landed on a head, you determined that Rory is not very good at copy-editing, dismissing all the evidence you have that she is, in fact, your peer. Consequently, you could remain steadfast in your belief about the number of mistakes in the text. Intuitively, your reaction would still be irrational despite your thinking that the Conciliatory View does not apply in this case.

Enoch then argues that this causes problems for the Conciliatory View. If we require agents to have justified beliefs about who their peers are, they need evidence for these beliefs. For example, in the Copy-Editing case, you and Rory have to be aware that you have a similar track record of being right about the number of mistakes in texts. In this case, as Enoch (2010) points out, the current disagreement must also be counted into the track record. But how? Equal Weight states that it could count either as your or as Rory's mistake, thus forcing you to give equal weight to both of your beliefs. In contrast, Enoch argues that in cases when you yourself are engaged in a disagreement with someone (in contrast to looking at someone else's disagreement), you can be justified to count this disagreement as the other person's mistake, thus demoting their epistemic status.

But why does Enoch think that? This strategy of using disagreements as evidence of other people's mistakes seems to beg the question, and Enoch acknowledges that. However, he thinks that it does it in a way that is not problematic: "The crucial point to note is that there is really nothing unique going on here. [W]hen deliberating epistemically about anything at all, your starting point is and cannot but be your own beliefs, degrees of beliefs, conditional probabilities, epistemic procedures and habits, and so on" (Enoch, 2010, p. 980). In other words, Enoch

thinks that the first-person perspective is ineliminable—the starting point of our epistemic endeavors is always our own epistemic states. As such, they *do* have a privileged position. In cases of peer disagreement, we are thus not considering two equal beliefs, for example, our belief that p and our peer’s belief $\neg p$, but the fact that p and our peer’s belief that $\neg p$. Consequently, we should count the belief that $\neg p$ as a mistake and reduce the epistemic status of our peer.⁸

Enoch thus rejects the Conciliatory View. However, he does not endorse the idea that we can simply remain steadfast in the face of peer disagreement. He thinks that we could react to disagreement “by simultaneously reducing [our] confidence in the controversial claim and in the reliability of both [us] and [our] (supposed) peer, though reducing it more sharply regarding [our] (supposed) peer” (Enoch, 2010, p. 993). That said, he does not argue for a general policy of determining these changes in specific cases. Rather, he states that the degree of confidence about the controversial belief you are justified to hold after learning about disagreement will depend on other factors, such as the “other things you (justifiably) believe, [other] evidence you have, [the] epistemic methods you are justified in employing” (Enoch, 2010, p. 994).

I do not find this argument entirely convincing. Regardless of the validity of Enoch’s claim about the priority of the first-person perspective, I think his argument engages in a kind of double-counting. Specifically, it seems that disagreement can act either as a reason to reduce confidence in the belief in question or as a reason to demote the epistemic status of the disagreeing peer. But it cannot act as both. For disagreement to act as a reason for reducing confidence in the belief in question, we must take it as evidence of our possible mistake. On the other hand, if we want to use it to demote our epistemic peer, we must see it as *our peer’s mistake*. However,

⁸Arguments along similar lines can be made based on the notion of epistemic self-trust. Epistemic self-trust is an attitude we can adopt towards our cognitive resources; it gives us a reason to remain steadfast in our beliefs despite the possibility of making mistakes (Foley, 2001). This argument can be extended to the cases of peer disagreement: when faced with peer disagreement, self-trust gives one a reason to privilege one’s own belief—the result of one’s cognitive resources which one trusts and can rely on—over the belief of the disagreeing peer (Rattan, 2014). However, as Rattan (2014) already points out, this argument misses a crucial distinction between the general threat of making a mistake and cases of peer disagreement, which can be seen as evidence that a mistake is not only possible but has indeed been made.

once we do that, the disagreement will already be accounted for.

Perhaps Enoch is assuming that the general threat of making a mistake is at issue. As he states, when you consider how to respond to your disagreement with Rory, you are not considering two equivalent beliefs but the fact that there are 76 mistakes in the text and Rory's belief that there are 70. Thus, you conclude that she has made a mistake. But it is nevertheless still possible that the mistake is yours—ultimately, you are not infallible. Thus, it also seems reasonable to reduce your confidence—that is, at least to some degree, since not trusting your epistemic faculties in this way would lead to skepticism. In this way, we can see how one disagreement can count both as evidence that perhaps we have made a mistake and that our peer *did* make a mistake simultaneously. However, as Rattan (2014) shows in his discussion of peer disagreement and epistemic self-trust, the general threat of making a mistake does not have the same conciliatory pressure as the evidence that an actual mistake has been made, which is provided by peer disagreement. Thus, when this evidence is accounted for by reducing the epistemic credentials of a peer, the remaining mere possibility that we are mistaken is not enough to prompt any significant belief revision.

This relates to Enoch's answer to the bootstrapping objection that can be found in Elga (2007) and was already presented in the previous section. Enoch (2010, p. 990) is aware of this objection but bites the bullet, contending that it is, "if not a particular instance then a close analogous of a very general worry [...] that we are not entitled to trust our own epistemic abilities to a degree greater than that which their track-record calls for [...]." Furthermore, he thinks that this worry quickly leads to skepticism. Between the two, he believes that skepticism is worse than the possibility of bootstrapping.

While Enoch's stance on bootstrapping might not be problematic as such, I think it presents a disadvantage in comparison to the Conciliatory View. To reiterate Rattan's point, the Conciliatory View is not inconsistent with epistemic self-trust. The Conciliatory View is a view about how to react to evidence of our mistakes. Self-trust, on the other hand, is about trusting our epistemic faculties, although they are fallible in general. In other words, it is consistent to (1) trust our epistemic faculties despite the general awareness that we sometimes make mistakes, that we

might be mistaken at this moment without recognizing it, and (2) at the same time recognize that we should also take seriously evidence that an actual mistake has been made. Since the Conciliatory View can fend off the worry of skepticism this way without biting the bullet about bootstrapping, it remains a promising view even in light of Enoch's considerations.⁹

Another argument against Equal Weight was put forward by Lasonen-Aarnio (2013). She argues that Equal Weight depends on an unmotivated constraint on our higher-order certainties, which she calls Indifference. Indifference states that if we understand higher-order conviction as a probability distribution over possible credences about p , Equal Weight implies that the two peers should assign equal probabilities to all possible credence about p even before learning about disagreement. More formally, Equal Weight can be interpreted as stating that $P_A(P(p) = r) = P_B(P(p) = r^*)$, where $P(p)$ is a first-order credence in p and higher-order conviction is represented as probability function over the first-order credence. Lasonen-Aarnio (2013) shows that accepting some minimal and uncontroversial assumptions about the situation of peer disagreement, it follows from Equal Weight that: $\forall r \forall r^* P_A(P(p) = r) = P_B(P(p) = r^*)$ (Indifference).

Indifference indeed seems to be a strong claim. It states that for peers to assign equal weight after learning about disagreement about p , they would already have had to assign equal weight to all possible beliefs about p beforehand. I agree with Lasonen-Aarnio that it is hard to motivate this claim. Reading in this way goes against the spirit of the Conciliatory View, which states that learning about peer disagreement should lower our higher-order conviction in the belief in question. Consequently, the above argument provides a strong reason to dismiss Equal Weight and, consequently, the Conciliatory View of peer disagreement.

There is another way of looking at Indifference that makes it much less threatening. Consider Lasonen Aarnio's articulation of Equal Weight again. It states that in cases of disagreement, peers should judge that it's equally likely that their first-order credences about p are correct, or $P_A(P(p) = r) = P_B(P(p) = r^*)$. Indifference then states that peers should judge all their possible credences about p in this way: for any credence $P(p) = r$, if a peer assumes a different credence, $P(p) = r^*$, these two

⁹I'm much indebted to the reviewer for helping me clarify the passages on Enoch (2010).

credences should be considered as equally likely. As pointed out above, this seems quite troublesome. But is it? Another way of looking at Indifference would be to accept it as a natural consequence of taking Equal Weight seriously. If we assume that peers should give equal weight to their opinions about p , it seems natural to think of this as a blanket requirement—it shouldn't depend on what credence about p they actually assume. However, that does not mean that should never assign more higher-order probability to one of these credences. Equal Weight only applies to cases of actual peer disagreement: if one assumes that it's highly likely that $P(p) = r$ and no one disagrees, one remains justified in this high higher-order probability pace Indifference.

This limitation on Indifference stems from the fact that merely possible disagreements do not have the same epistemic value as real ones. This has already been extensively discussed in the literature. Christensen (2007), for example, argues that actual disagreements are informative because they give us some evidence that this possibility of error has actualized. Merely possible disagreement, on the other hand, gives us no new such evidence.

Carey (2011) makes a very similar argument, only in a more precise way. Consider the Copy-Editing case again. Let's assume that when you and Rory copy-edit the same texts, you agree about the number of mistakes in them and are right about that 80% of the time; 10% of the time, you agree but are both incorrect; 10% of the time, you disagree—half of that time you are right and Rory is incorrect, and half of that time it's the other way around. Now, when you first form a belief about the number of mistakes in the text, you can be confident with 0.8 credence that you are right. But then you learn that Rory disagrees with you. This means that you have found yourself in that 10% of cases where you and Rory disagree. You know that, in such cases, it is equally likely that either you or Rory is right. Thus, you should give your beliefs equal weight.

Now, we can compare this with merely possible disagreements. It is *always* possible that you and Rory would disagree. The probability of a mere possible disagreement between you and Rory is thus 1. But if that is so, then merely possible disagreements give us no new information and cannot act as higher-order evidence.

Because merely possible disagreements do not act as higher-order evidence, we

must restrict Indifference to cases of actual peer disagreement. But if we do that, it loses much of its strength—it becomes a natural and intuitive consequence of Equal Weight. One might still worry that it is too strong since it gives unintuitive answers in cases such as the Black Death. If Rory and Max are peers and assume different credences to the proposition that the medieval plague pandemic happened, they should consider their credence as equally likely, as per Indifference. However, as we saw in the previous section, Rory has other independent reasons for dismissing Max’s opinion as a genuine case of peer disagreement. Thus, I think we don’t need to reject Equal Weight based on the argument provided by Lasonen-Aarnio (2013).

5 Defending First-Order Import

I will now discuss the claim that evidence of our unreliability about p (for example, the fact of peer disagreement) should prompt us to change our belief about p . This claim is necessary if we want to argue that disagreeing peers should change their *beliefs* after learning about disagreement. In other words, First-Order Import provides a bridge between higher-order uncertainty and first-order beliefs.

First-Order Import, or the idea that learning about peer disagreement (or receiving other kinds of higher-order evidence) should prompt us to revise our beliefs, has received a lot of criticism. Some authors have argued that such evidence does not have first-order import (Kelly, 2005; Lasonen-Aarnio, 2014; Littlejohn, 2018; Titelbaum, 2015). Others have argued that conciliating in the face of peer disagreement requires us to ignore evidence (Kelly, 2010; Sliwa and Horowitz, 2015), leads to unwanted skepticism (see Frances and Mathenson, 2018), violates coherence of beliefs (Jehle and Fitelson, 2009), makes rationality come by too easily (Kelly, 2010), leads to epistemic akrasia (Christensen, 2020), is rationally toxic (Christensen, 2016), and is self-defeating (Dixon, 2024).

I will argue that none of these arguments directly target the Conciliatory View. Instead, they all concern a more general issue of whether and how we should consider higher-order evidence. As I see it, the above arguments fall into two distinct groups. The first group of arguments rejects the general idea that higher-order evidence should impact our beliefs, thus rejecting the view that we should

conciliate with a disagreeing peer. The arguments in the other group assume that the Conciliatory View relies on a specific interpretation of the First-Order Import and argue against that interpretation.

So, instead of directly addressing these arguments, I will contend myself with showing that (1) First-Order Import is not implausible given the current state of the literature on higher-order evidence, and (2) that the effectiveness of some arguments against the Conciliatory View depends on the specific interpretation of First-Order Import. I will particularly focus on arguing that the Conciliatory View can avoid the charge of self-defeat. Keeping with the spirit of the main claim of this paper—that we should keep the hope for the Conciliatory View alive—I will not attempt to show that First-Order Import is certainly true. Rather, my goal will be to show that, when understood correctly, it is a plausible principle about higher-order evidence.

Let's first look at the argument about the plausibility of First-Order Import. As mentioned, First-Order Import concerns the relationship between first- and higher-order convictions; it answers whether and how higher-order evidence should affect our beliefs. Three general responses to this question have been suggested in the literature (see Horowitz, 2022 and Ye, 2022 for recent reviews). One response, sometimes called level-splitting, steadfastness, or the limits of defeat view, denies the import of higher-order evidence on our beliefs. This response is motivated by the analysis that higher-order evidence generates puzzles that are hard to accommodate within a consistent account of rationality.

In contrast, calibrationism, or the higher-order defeat view, states that higher-order evidence should have some impact on our beliefs, either defeating the relation between our beliefs and the first-order evidence (in the case of negative higher-order evidence) or strengthening it (in the case of positive evidence). This view is motivated by our intuitive judgments in cases such as Copy-Editing. In addition, proponents of the view suggest that violating it would make us guilty of bootstrapping and dogmatic reasoning (Elga, 2007) and possibly make us epistemically akratic (Henderson, 2022).

The third response to the question of higher-order evidence, sometimes called the dilemma or two-norms view, combines the previous two answers. It states that higher-order evidence presents genuine rational dilemmas where all available

responses are problematic. On the one hand, it concedes that higher-order evidence creates puzzles that are hard to account for in a consistent account of epistemic rationality. On the other hand, it argues that given our less-than-ideal epistemic situation, we should nevertheless act upon higher-order evidence.

First-Order Import most naturally follows from calibrationism or the higher-order defeat view but is also consistent with some dilemma views of higher-order evidence. However, it is denied by the level-splitting or the steadfast view. If the latter view was the dominant answer to the question of higher-order evidence, we would thus have a good reason to reject First-Order Import. That said, this does not seem to be the case. That is not to say that the steadfast or the level-splitting view does have its share of supporters, among which Kelly (2005) and Titelbaum (2015) explicitly use it to argue against the Conciliatory View of peer disagreement (for answers, see Field, 2019; Kelly, 2010), but it is by no account a dominant view.

The question of whether higher-order evidence has first-order import goes beyond well beyond the discussion of peer disagreement. It is a topic of lively debate in contemporary epistemology, and it would be futile for me to attempt an answer here. In any event, the debate about the epistemic importance of higher-order evidence and, consequently, about principles such as First-Order Import is still open. In addition, First-Order Import is consistent with intuitive answers in cases like Hypoxia and Copy-Editing. It is also in some form assumed by other views on peer disagreement that are otherwise incompatible with the Conciliatory View but suggest that we must at least sometimes revise our beliefs in the face of peer disagreement, such as the Justificatory View (Lackey, 2010) and the Total Evidence View (Kelly, 2013). While it would be premature to accept it as the correct solution to the question of higher-order evidence, it is most definitely a plausible and reasonable working hypothesis (for a recent defense, see Ye (2022)).

Next, consider the argument that the effectiveness of some arguments against the First-Order Import relies on authors assuming that the Conciliatory View adopts a specific interpretation of the claim. To see the dynamic of this argument at play, I will first look at the critique by Jehle and Fitelson (2009). Then, I will show how this strategy can be used to reject a much more pressing charge that the Conciliatory View is almost certainly self-undermining (Dixon, 2024).

Jehle and Fitelson (2009) ask what it means to give equal weight to peers in cases of disagreement. They propose an interpretation of First-Order Import under which degrees of belief directly respond to changes in higher-order convictions. Specifically, under the interpretation they call Straight Averaging, peers' new credences (P^1) after learning about the disagreement should equal $P_A^1(p) = P_B^1(p) = \frac{P_A^0(p) + P_B^0(p)}{2}$.¹⁰ In other words, Straight Averaging states that the disagreeing peers should resolve their disagreement by adopting an average of their prior credences as their new credences. Jehle and Fitelson (2009) then argue that Straight Averaging fails as a rational principle since it can sometimes force reasoners to adopt an incoherent set of credences.

To see what Jehle and Fitelson mean, consider this example. Let's say that you and Rory are peers with regard to propositions $p \wedge q$ and $\neg p \wedge q$, but not about propositions $p \wedge \neg q$ and $\neg p \wedge \neg q$.¹¹ Imagine that you disagree about the peer propositions; after learning about this, you use Straight Averaging to update your credences about them. Now, here comes the problem. Bayesian epistemology requires that credences are probability functions. This implies that an agent's credences about a set of mutually exclusive and collectively exhaustive propositions adds up to 1. Let's assume that this was the case for your and Rory's credences before you learned about the disagreement. But to make sure that the requirement is still satisfied after you apply Straight Averaging to your credences about peer propositions, you must also change your credences about related non-peer propositions. Unfortunately, Straight Averaging does not speak to that. Therefore, it's either incomplete, or, if we read it as a requirement that credences about non-peer propositions *don't* change, it can cause agents to have incoherent credences.

However, it is essential to notice that this argument is limited in scope: it concerns only a specific interpretation of First-Order Import, not the claim itself. The

¹⁰Note that Jehle and Fitelson (2009) take Straight Averaging as an articulation of the Conciliatory View (which they call "The Equal Weight View"). Thus, Straight Averaging is more substantial than an interpretation of First-Order Import should be since it already assumes Equal Weight. A properly limited reading would say something like: Evidence of peer disagreement should prompt one to adopt a new credence, calculated as a weighted average of peers' original credences. Together with Equal Weight, this principle yields Straight Averaging.

¹¹Is a situation like this possible? Jehle and Fitelson (2009) think so and present a convincing example in footnote 9 of their paper.

problem can thus be responded to by providing an alternative interpretation of First-Order Import. This strategy is already employed by Jehle and Fitelson (2009), so I will briefly outline it here. They suggest that we should modify Straight Averaging with a clause about making minimal changes to non-peer propositions to satisfy the coherence requirement. The upshot of their discussion is thus that there are some difficulties and potential trade-offs involved in providing a Bayesian interpretation of First-Order Import and, consequently, the Conciliatory View. Crucially for us, however, they do not show that First-Order Import is untenable as such.

This strategy of providing a more precise articulation of First-Order Import can also be used to avoid the charge that the Conciliatory View is almost certainly self-undermining, presented recently by Dixon (2024). This charge is quite serious: Dixon claims that under any interpretation of First-Order Import that gives intuitive results in cases such as Copy-Editing, the Conciliatory View ends up self-defeating. He presents the following deceptively simple but strong argument against the Conciliatory View:

1. “If acknowledged peers recognize they disagree on p , then peers should suspend belief on p . (Conciliationism)
2. Acknowledged peers recognize they disagree on Conciliationism. (Empirical premise)
3. Thus, peers should suspend belief in Conciliationism.” (Dixon, 2024, p. 3)

Dixon’s argument targets a specific belief-suspension interpretation of First-Order Import, which states that evidence of peer disagreement about p defeats the reasons one has for their belief about p . He considers this to be “the most well-motivated version” of the Conciliatory View (Dixon, 2024, p. 4). Furthermore, he argues that even if we disagree with this and prefer a different interpretation of the Conciliatory View, this almost certainly will not allow us to defend the view against his challenge. In what remains of this section, I will first motivate a move away from Dixon’s belief-suspension picture of the Conciliatory View and toward a degrees-of-belief framework. Then, I will show how we can resist the argument in this framework.

The primary motivation for moving away from the general suspension of belief reading of First-Order Import is its limited scope. It only applies to cases in which one peer believes p and the other $\neg p$. It does not give clear prescriptions where one of the peers already suspends judgment about p . Intuitively, we would also like to say something about such cases. For example, consider two scientists who disagree whether treatment T has effect p . One scientist thinks there is no effect, while the other judges that there is not enough evidence to make this call. In other words, she thinks that the evidence could still go either way and thus suspends judgment about the question. The two scientists disagree, but it is not clear how Dixon's "Conciliationism" applies to their case. Meanwhile, a credal version of the Conciliatory View, e.g., Straight Averaging, would have no problem with providing an answer in this case.

What Dixon calls Conciliationism is thus a limited interpretation of First-Order Import. Thus, we can reject his argument by rejecting the first premise and providing a different interpretation of the view. However, Dixon already anticipates this response and provides an answer to it. He gives this argument:

1. "If argument A adequately defends the Conciliatory View from the self-undermining challenge, then epistemic peers do not disagree with the reasoning or auxiliary premises of A.
2. Epistemic peers (almost certainly) disagree with the reasoning or auxiliary premises of A.
3. Thus, A (almost certainly) fails to adequately defend the Conciliatory View from the self-undermining challenge." (Dixon, 2024, p. 8)

This argument extends the logic of the initial argument against the Conciliatory View to all possible defenses of the view: if peers disagree about a premise used in the defense of the Conciliatory View, then this premise is, by the view's own lights, undermined. For example, let's say we argue that a specific interpretation of First-Order Import, call it Import*, defends the Conciliatory View against Dixon's self-undermining argument. But assume that Import* itself is not uncontroversial (which is highly plausible given the state of the literature). If that is the case, we

should, in Import*'s own light, conciliate with our peers about Import*. But if we should conciliate about Import*, we cannot use Import* as a justified premise in our argument in defense of the Conciliatory View.

This is a persuasive argument. But I do not think it is correct. Specifically, premise (1) does not seem to me to necessarily follow from the Conciliatory View. In other words, we can construct examples where agents correctly follow a credal version of the view but are not susceptible to the above argument. We can already get one such example by modifying the case of the two scientists from above. Two philosophers disagree on whether Straight Averaging is a correct interpretation of the First-Order Import: one thinks that it is, while the other suspends judgment about it. In this case, the principle would not be self-undermining: by applying Straight Averaging, the first philosopher could still retain a relatively high credence in it.

Notice that this only helps in a very specific set of cases. To exclude them, we could amend the second premise of the above argument to read: (2*) "Epistemic peers (almost certainly) disagree with the reasoning or auxiliary premises of A, such that some peers believe A while others believe not-A." Although (2*) is stronger than (2), I still find it convincing, especially in the context of a philosophical debate.

However, even with (2*), it is still possible to construct counter-examples to Dixon's argument once we are considering credal versions of the Conciliatory View. Let's say we think the weakened version of Straight Averaging, Straight Averaging*, which says that our new post-disagreement credences should be calculated as a weighted average of peers' original credences, is the right interpretation of First-Order Import. Together with the other claims made by the Conciliatory View, Straight Averaging* does not imply what Dixon calls Conciliationism: "If acknowledged peers recognize that they disagree on p , then peers should suspend belief on p ." Consequently, we have an answer to Dixon's self-undermining challenge. Now, this is where the second part of Dixon's argument kicks in: there will, of course, be disagreement about this, so sooner or later, we should also conciliate about, and thus undermine, Straight Averaging*. But consider one specific such scenario: let's say that most of the disagreeing peers are quite positive about the view (although with a range of different credences). At the same time, there are only

a few strong dissenters. In such a case, it is not unimaginable that after conciliating, we end up with a relatively high credence in the view. Perhaps even a credence that would warrant the belief that Straight Averaging* is the right view to have about First-Order Import.

The upshot of this case is that we can still believe that First-Order Import should be interpreted as Straight Averaging* despite there being substantial disagreement about it. Thus, it is possible, pace premise (1) from above, to imagine that an argument adequately defends the Conciliatory View, despite the fact that epistemic peers *do* disagree with relation to this argument. This is then a counter-example to Dixon's self-undermining challenge.

To elaborate, I do not wish to endorse Straight Averaging* as a correct interpretation of First-Order Import. All I want to show is that it is possible to develop a credal version of the Conciliatory View that can, in principle, walk the goldilocks line between being strong enough to give intuitive answers in cases such as Copy-Editing while not being self-undermining. Dixon claims that any such attempt will almost certainly fail. But the above case does not seem especially implausible or contrived and cannot, in my mind, be dismissed on these grounds. Thus, this keeps the hope alive, at least for a credential version of the Conciliatory View.¹²

6 Concluding Defense

The previous three sections of this paper defended different aspects of the Conciliatory View of peer disagreement. This concluding section will defend the general approach taken up in the paper. I will defend my approach against two charges: (1) it amounts to misguided gerrymandering of epistemological principles, and (2) it ends up with a view that does not resemble what is commonly understood as the Conciliatory View.

More specifically, the gerrymandering objection states that my defense of the Conciliatory View consists of assigning only unproblematic claims to the view while unloading the more troublesome ones onto the higher-order evidence debate.

¹²I must again thank the reviewer for helping me to clarify the discussion of the self-undermining worry.

To recall, the primary strategy was to separate the claims made by the view specific to peer disagreement debate and claims that concern higher-order evidence more generally. However, I think this was done with some care. For example, although Independence concerns the problem of higher-order evidence more generally, I provided a novel defense of the principle. Some might nonetheless find it problematic that the exact meaning of First-Order Import was left unspecified. However, I see this as an advantage of my approach rather than a downside. It shows that the Conciliatory View is compatible with multiple accounts of how higher-order evidence affects our beliefs.

The second worry asks: is the view defended here—with so much still left open—still what is typically understood as the Conciliatory View? For example, what if it turns out that First-Order Import is correct, but only in a very minimal sense? In other words, what if it turns out that even evidence of our complete unreliability about p should only have a very minimal effect on our beliefs? Would such a minimal reading get the Conciliatory View of peer disagreement off the ground?

Probably not. Such a view would not give intuitively correct answers in cases like Copy-Editing. However, I do not find this option very likely. Since cases of peer disagreement are also cases of higher-order evidence, every intuitive account of the epistemic role of higher-order evidence would already have to consider these cases. Thus, if it turns out that higher-order evidence does have first-order import, the Conciliatory View of peer disagreement will likely follow.

Acknowledgments

Some of the research for this paper was conducted as part of my master's thesis at the University of Ljubljana, Faculty of Arts. I would foremost like to thank my supervisors, Olga Markič, Dunja Šešelja, and Borut Trpin, for their selfless help and insightful comments. I'm also grateful to the anonymous reviewer of *Episteme*—their comments made this a much better paper.

References

- Bogardus, Tomas (2009). “A Vindication of the Equal-Weight View”. In: *Episteme* 6 (3), pp. 324–335. DOI: 10.3366/E1742360009000744.
- Carey, Brandon (2011). “Possible disagreements and defeat”. In: *Philosophical Studies* 155 (3), pp. 371–381. DOI: 10.1007/s11098-010-9581-5.
- Christensen, David (2007). “Epistemology of Disagreement: The Good News”. In: *The Philosophical Review* 116 (2), pp. 187–217. DOI: 10.1215/00318108-2006-035.
- (2010). “Higher-order evidence”. In: *Philosophy and Phenomenological Research* 81 (1), pp. 185–215. DOI: 10.1111/j.1933-1592.2010.00366.x.
- (2011). “Disagreement, Question-Begging and Epistemic Self-Criticism”. In: *Philosophers’ Imprint* 11 (6).
- (2016). “Conciliation, Uniqueness and Rational Toxicity”. In: *Nous* 50 (3), pp. 584–603. DOI: 10.1111/nous.12077.
- (2019). “Formulating Independence”. In: *Higher-Order Evidence*, pp. 13–34. DOI: 10.1093/OSO/9780198829775.003.0001.
- (2020). “Akratic (epistemic) modesty”. In: *Philosophical Studies* 178 (7), pp. 2191–2214. DOI: 10.1007/S11098-020-01536-6.
- Cocchiaro, Mariangela Zoe and Bryan Frances (2021). “Epistemically Different Epistemic Peers”. In: *Topoi* 40 (5), pp. 1063–1073. DOI: 10.1007/S11245-019-09678-X/TABLES/3.
- Cruz, Helen De and Johan De Smedt (2013). “The value of epistemic disagreement in scientific practice. The case of Homo floresiensis”. In: *Studies in History and Philosophy of Science Part A* 44 (2), pp. 169–177. DOI: 10.1016/j.shpsa.2013.02.002.
- Dixon, Jonathan (2024). “No hope for conciliationism”. In: *Synthese* 203 (5), p. 148. DOI: 10.1007/s11229-024-04590-3.
- Dorst, Kevin (2020). “Evidence: A Guide for the Uncertain”. In: *Philosophy and Phenomenological Research* 100 (3), pp. 586–632. DOI: 10.1111/PHPR.12561.

- Douven, Igor (2010). “Simulating peer disagreements”. In: *Studies in History and Philosophy of Science Part A* 41 (2), pp. 148–157. DOI: 10.1016/j.shpsa.2010.03.010.
- Elga, Adam (2007). “Reflection and Disagreement”. In: *Nous* 41 (3), pp. 478–502. DOI: 10.1111/j.1468-0068.2007.00656.x.
- (2008). “Lucky to be rational”.
- (2010). “How to Disagree About How to Disagree”. In: ed. by Richard Feldman and Ted A. Warfield. Oxford University Press. DOI: 10.1093/acprof:oso/9780199226078.001.0001.
- Enoch, David (2010). “Not Just a Truthometer: Taking Oneself Seriously (but not Too Seriously) in Cases of Peer Disagreement”. In: *Mind* 119 (476), pp. 953–997. DOI: 10.1093/mind/fzq070.
- Feldman, Richard (2005). “Respecting the Evidence”. In: *Philosophical Perspectives* 19 (1), pp. 95–119. DOI: 10.1111/j.1520-8583.2005.00055.x.
- (2006). “Epistemological Puzzles about Disagreement”. In: Oxford University Press Oxford, pp. 199–215. DOI: 10.1093/oso/9780199273317.003.0013.
- (2009). “Evidentialism, Higher-Order Evidence, and Disagreement”. In: *Episteme* 6 (3), pp. 294–312. DOI: 10.3366/E1742360009000720.
- Field, Claire (2019). “It’s OK to Make Mistakes: Against the Fixed Point Thesis”. In: *Episteme* 16 (2), pp. 175–185. DOI: 10.1017/EPI.2017.33.
- Foley, Richard (2001). *Intellectual Trust in Oneself and Others*. New York: Cambridge University Press.
- Frances, Bryan and Jonathan Mathenson (2018). “Disagreement”. In:
- Henderson, Leah (2022). “Higher-order evidence and losing one’s conviction”. In: *Noûs* 56 (3), pp. 513–529. DOI: 10.1111/nous.12367.
- Horowitz, Sophie (2022). “Higher-Order Evidence”. In: Fall 2022. Metaphysics Research Lab, Stanford University.
- Jehle, David and Branden Fitelson (2009). “What is the “Equal Weight View?”” In: *Episteme* 6 (3), pp. 280–293. DOI: 10.3366/E1742360009000719.
- Kelly, Thomas (2005). “The Epistemic Significance of Disagreement”. In: ed. by Tamar Szabo Gendler and John Hawthorne.

- Kelly, Thomas (2010). “Peer Disagreement and Higher-Order Evidence”. In: Oxford University Press, pp. 111–174. DOI: 10.1093/acprof:oso/9780199226078.003.0007.
- (2013). “Disagreement and the Burdens of Judgment”. In: ed. by David Christensen and Jennifer Lackey. Oxford University Press, pp. 31–53. DOI: 10.1093/acprof:oso/9780199698370.003.0003.
- King, Nathan L. (2012). “Disagreement: What’s the Problem? or A Good Peer is Hard to Find”. In: *Philosophy and Phenomenological Research* 85 (2), pp. 249–272. DOI: 10.1111/J.1933-1592.2010.00441.X.
- Kornblith, Hilary (2010). “Belief in the Face of Controversy”. In: Oxford University Press, pp. 29–52. DOI: 10.1093/acprof:oso/9780199226078.003.0003.
- Lackey, Jennifer (2010). “A Justificationist View of Disagreement’s Epistemic Significance”. In: Oxford University Press, pp. 298–325. DOI: 10.1093/acprof:oso/9780199577477.003.0015.
- Lasonen-Aarnio, Maria (2013). “Disagreement and Evidential Attenuation”. In: *Noûs* 47 (4), pp. 767–794. DOI: 10.1111/NOUS.12050.
- (2014). “Higher-Order Evidence and the Limits of Defeat”. In: *Philosophy and Phenomenological Research* 88 (2), pp. 314–345. DOI: 10.1111/PHPR.12090.
- Littlejohn, Clayton (2018). “Stop Making Sense? On a Puzzle about Rationality”. In: *Philosophy and Phenomenological Research* 96 (2), pp. 257–272. DOI: 10.1111/PHPR.12271.
- Lord, Errol (2014). “From Independence to Conciliationism: An Obituary”. In: *Australasian Journal of Philosophy* 92 (2), pp. 365–377. DOI: 10.1080/00048402.2013.829506.
- Matheson, Jonathan (2009). “Conciliatory Views of Disagreement and Higher-Order Evidence”. In: *Episteme* 6 (3), pp. 269–279. DOI: 10.3366/E1742360009000707.
- (2015). “Are Conciliatory Views of Disagreement Self-Defeating?” In: *Social Epistemology* 29 (2), pp. 145–159. DOI: 10.1080/02691728.2014.907833.
- (2024). “Disagreement and Higher-Order Evidence”. In: *The Routledge Handbook of the Philosophy of Evidence*. Ed. by Maria Lasonen-Aarnio and Clayton Littlejohn. Routledge.

- Moon, Andrew (2018). “Independence and New Ways to Remain Steadfast in the Face of Disagreement”. In: *Episteme* 15 (1), pp. 65–79. DOI: 10.1017/epi.2016.42.
- Rasmussen, Mattias Skipper, Asbjorn Steglich-Petersen, and Jens Christian Bjerring (2018). “A higher-order approach to disagreement”. In: *Episteme* 15 (1), pp. 80–100. DOI: 10.1017/EPI.2016.43.
- Rattan, Gurpreet (2014). “Disagreement and the First-Person Perspective”. In: *Analytic Philosophy* 55 (1), pp. 31–53. DOI: 10.1111/phib.12038.
- Sliwa, Paulina and Sophie Horowitz (2015). “Respecting all the evidence”. In: *Philosophical Studies* 172 (11), pp. 2835–2858. DOI: 10.1007/s11098-015-0446-9.
- Titelbaum, Michael G. (2015). “Rationality’s Fixed Point (or: In Defense of Right Reason)”. In: Oxford University Press, pp. 253–294. DOI: 10.1093/acprof:oso/9780198722762.003.0009.
- Ye, Ru (2022). “Higher-Order Evidence and Calibrationism”. In: *Elements in Epistemology*. DOI: 10.1017/9781009127332.