



Marc Andree Weber*

Conciliatory Views on Peer Disagreement and the Order of Evidence Acquisition

<https://doi.org/10.1515/krt-2021-0023>

Published online September 29, 2021

Abstract: The evidence that we get from peer disagreement is especially problematic from a Bayesian point of view since the belief revision caused by a piece of such evidence cannot be modelled along the lines of Bayesian conditionalisation. This paper explains how exactly this problem arises, what features of peer disagreements are responsible for it, and what lessons should be drawn for both the analysis of peer disagreements and Bayesian conditionalisation as a model of evidence acquisition. In particular, it is pointed out that the same characteristic of evidence from disagreement that explains the problems with Bayesian conditionalisation also suggests an interpretation of suspension of belief in terms of imprecise probabilities.

Keywords: disagreement, peers, equal weight view, Bayesian conditionalisation, evidence, imprecise probabilities

1 Introduction

Let us assume that Richard's credence towards the proposition that Pete will drink more than three beers tonight is $\frac{1}{4}$, while Siena's credence towards the same proposition is $\frac{3}{4}$. Richard and Siena know Pete's drinking behaviour equally well, and they know from past experience that they are equally good at predicting how many beers he will have. When they then exchange views, Richard gets evidence that he has underestimated Pete's thirst for beer, while Siena gets evidence that she has overestimated it. Hence, it appears natural that Richard should raise his credence and Siena lower hers.

Richard and Siena are *epistemic peers* concerning Pete's drinking behaviour: they are equally competent and knowledgeable with regard to predicting it.¹

¹ There are other and, in my opinion, better ways to define epistemic peerhood – see, e.g., my Weber (2017), in which I defend a likelihood account of epistemic peerhood – but they do not matter for our purposes, so let's stick with this rough but simple characterisation.

*Corresponding author: Marc Andree Weber, Philosophisches Seminar, University of Mannheim, L 9, 5, 68161 Mannheim, Germany, E-mail: marc.andree.weber@phil.uni-mannheim.de

So-called *conciliatory views* hold that they should revise their original credences at least a bit as soon as they become aware of their peer disagreement. An important special case of a conciliatory view is the *Equal Weight View* (henceforth *EW*), which holds that they should give the credences of their respective epistemic peer the same weight that they give their own.

Conciliatory views are often said to apply, not only when it comes to predicting specific aspects of human behaviour, but also in those peer disagreement cases in which we seem unable to find out who is right. Thus, if Richard and Siena are epistemic peers concerning 19th-century Russian literature and Richard thinks that Gogol is a greater writer than Dostoevsky, while Siena does not, they should revise their respective credences; if they are epistemic peers concerning political matters and Siena thinks that their country should impose sanctions against Iran, while Richard does not, they should revise their respective credences; and if they are epistemic peers concerning ethics and Richard thinks that average utilitarianism is the best moral theory, while Siena does not, they should revise their respective credences as well. Fortunately, we need not care here whether conciliatory views cover those kinds of peer disagreements, too; my concern in this paper is their *reconciliation with Bayesianism*.

In more detail: In Section 2, I will point out that the standard interpretation of conciliatory views is incompatible with Bayesian conditionalisation because the order in which one acquires new evidence matters for the former but not for the latter. In Section 3, I will argue that a specific feature of evidence from disagreement, its so-called *retrospective aspect*, suggests that a particular order of evidence acquisition is preferable in many cases, and will indicate which cases are exceptions. Finally, in Section 4, I will present an alternative interpretation of EW that is more in line with the retrospective aspect, and will explore this interpretation's consequences for updating beliefs in a broadly Bayesian way.

In order to tackle these issues, using credence talk is helpful but ultimately inessential. It is helpful insofar as it simplifies the presentation a lot; it is inessential insofar as all that follows could be reformulated in terms of just three doxastic attitudes – belief, disbelief, and suspension of belief – instead of continuum many. We could, for example, say that conciliatory views require epistemic peers to give at least *some* weight to the others' beliefs, thereby leaving it open whether a specific disagreement would call for a change of doxastic attitude. I will take up this point towards the end of the paper.

It also simplifies matters if we focus, not on conciliatory views in general, but on EW. At least the results in Sections 2 and 3 could be reformulated such that they apply to all conciliatory views.

2 Splitting the Difference and Bayesian Conditionalisation

According to the standard interpretation of EW, ‘give the credences of your epistemic peers the same weight that you give your own’ just means that you should adopt a credence that equals the arithmetic mean of your own and your epistemic peers’ original credences. This is called *splitting the difference*. Let us assume in this and the following section that *splitting the difference* is the correct interpretation of EW. So, if $C_R(h)$ is Richard’s credence towards hypothesis h and $C_S(h)$ Siena’s, *splitting the difference* requires them to revise their beliefs in the case of a peer disagreement D as follows:

$$C_R^{+D}(h) = C_S^{+D}(h) = \frac{1}{2} C_R(h) + \frac{1}{2} C_S(h). \quad (1)$$

If we apply *splitting the difference* to the example with which we began, both Richard and Siena should adopt credence $\frac{1}{2}$ towards the proposition that Pete will drink more than three beers tonight.

More generally, conciliatory views require Richard to adopt

$$C_R^{+D}(h) = x C_S(h) + (1-x)C_R(h), \quad \text{for some } x \text{ with } 0 < x \leq \frac{1}{2}, \quad (2)$$

and Siena to adopt

$$C_S^{+D}(h) = x C_R(h) + (1-x)C_S(h), \quad \text{for some } x \text{ with } 0 < x \leq \frac{1}{2}. \quad (3)$$

Now, assume that Richard and Siena receive new evidence e for assessing h . Then, according to Bayesian conditionalisation:

$$C_{R/S}^{+e}(h) = C_{R/S}(h|e) =_{\text{def}} \frac{C_{R/S}(h \wedge e)}{C_{R/S}(e)}, \quad (4)$$

where the subscript R/S means that the respective credence is both Richard’s and Siena’s. The second equation in (4) is simply the definition of conditional probability.

Example 1. Assume that $C_R(e) = \frac{1}{2}$, $C_R(h \wedge e) = \frac{1}{4}$, $C_S(e) = \frac{1}{3}$ and $C_S(h \wedge e) = \frac{1}{4}$. What should Richard and Siena believe if they come to know both that e holds and that they disagree about h ? We have

$$C_R^{+e}(h) = C_R(h|e) = \frac{C_R(h \wedge e)}{C_R(e)} = \frac{\frac{1}{4}}{\frac{1}{2}} = \frac{1}{2}, \quad (5)$$

$$C_S^{+e}(h) = C_S(h|e) = \frac{C_S(h \wedge e)}{C_S(e)} = \frac{\frac{1}{4}}{\frac{1}{3}} = \frac{3}{4}, \quad (6)$$

$$C_{R/S}^{+e+D}(h) = \frac{1}{2}C_R^{+e}(h) + \frac{1}{2}C_S^{+e}(h) = \frac{1}{4} + \frac{3}{8} = \frac{5}{8}. \quad (7)$$

In (5) and (6), an instance of the conditionalisation rule, namely (4), has been applied, and in (7), we have used *splitting the difference*, as in (1).² If we proceed the other way around, that is, if we first split the difference twice and apply (4) only afterwards, we get

$$C_{R/S}^{+D}(h \wedge e) = \frac{1}{2}C_R(h \wedge e) + \frac{1}{2}C_S(h \wedge e) = \frac{1}{8} + \frac{1}{8} = \frac{1}{4}, \quad (8)$$

$$C_{R/S}^{+D}(e) = \frac{1}{2}C_R(e) + \frac{1}{2}C_S(e) = \frac{1}{4} + \frac{1}{6} = \frac{5}{12}, \quad (9)$$

$$C_{R/S}^{+D+e}(h) = C_{R/S}^{+D}(h|e) = \frac{C_{R/S}^{+D}(h \wedge e)}{C_{R/S}^{+D}(e)} = \frac{\frac{1}{4}}{\frac{5}{12}} = \frac{3}{5}. \quad (10)$$

Since $\frac{5}{8} \neq \frac{3}{5}$, we get $C_{R/S}^{+e+D}(h) \neq C_{R/S}^{+D}(h|e)$ – hence, Bayesian conditionalisation is violated. In other words: the kind of belief revision that is suggested by *splitting the difference* cannot be modeled in the usual Bayesian way.³ The same problem arises if we focus not on EW but on conciliatory views in general, that is, if we use (2) and (3) instead of (1).

EXAMPLE 1 shows that, given *splitting the difference*, it matters whether one first gets higher-order evidence from disagreement (D) or a piece of non-disagreement

2 In order to make the example work, we have to assume that e does not affect Richard and Siena's status as epistemic peers. Otherwise, after conditioning on e there would be no reason any more to split the difference, as in (7). The assumption, however, is not very restrictive: most of the time, the disagreements we have and the evidence that bears on them do not concern the reliability of the disagreeing parties. For instance, in the murder case that is described on p. 5, e does clearly not call Richard and Siena's peerhood status into question. (I am grateful to an anonymous referee for bringing to my attention the possibility that e affects peerhood statuses.)

3 Cf. Shogenji (2007) and Jehle and Fitelson (2009) as well as my Weber (2019, pp. 148–55). See also Wagner (1985, pp. 99–104), for a discussion in the earlier literature of the fact that arithmetic averaging and Bayesian conditionalisation do not commute.

evidence (e). It does make a difference whether two peers first acquire, separately from each other, some additional piece of first-order evidence and then come to know what the other one thinks of the totality of the evidence, or whether they first come to know what the other one thinks of the totality of the evidence minus that additional piece and then acquire, separately from each other, the additional piece of evidence. In other words, given *splitting the difference*, the order of evidence acquisition matters for what we should believe.

This contradicts the Bayesian model of belief revision, since within this model, it does not matter which piece of evidence one gets first – and, some would add, this is an intuitively correct understanding of how we process evidence. Take, as an example, a criminal case. It appears to be irrelevant whether the detectives first interview witness A and then witness B , or vice versa.

Our problem is twofold. First, since *splitting the difference* and the Bayesian framework are not compatible, the question is which has to go.⁴ (Surprisingly, my answer will be: both.)

Second, since we get different results for different orders of evidence acquisition, the question arises which order of evidence acquisition is better, epistemically speaking. Assume, for example, that Richard and Siena are detectives who investigate, independently from each other, in a murder case. Let h be the proposition that the butler is the murderer, and e the proposition that the gardener saw the butler in the billiard room just before the crime happened. Should Richard and Siena *exchange views first*, that is, first discuss matters and afterwards interview the gardener separately in order to find out whether e holds, or should they *exchange views last*, that is, first incorporate the evidence that they get from talking to the gardener, and then discuss matters?⁵ Put differently: If we assume again that $C_R(e) = \frac{1}{2}$, $C_R(h \wedge e) = \frac{1}{4}$, $C_S(e) = \frac{1}{3}$, and $C_S(h \wedge e) = \frac{1}{4}$, should they proceed as in (5)–(7), or should they proceed as in (8)–(10)? (My answer will be: first take in new evidence, then exchange views. Richard and Siena are well advised to separately incorporate the evidence they get from interviewing the gardener before talking to

4 As Shogenji shows, substituting the arithmetic mean value by a specifically weighted arithmetic average or by other mean values such as the geometric one is of no help; no mean value satisfies the standard requirements of the Bayesian model (Shogenji 2007, pp. 8–12). Jehle and Fitelson discuss several alternatives to averaging (Jehle and Fitelson 2009, pp. 287–90), but none of them seems to capture the spirit of EW. See also Moss (2012) and Staffel (2015) for compromising strategies, which require agents to maximise, for instance, average expected epistemic utility, as alternatives to splitting the difference.

5 In real-world cases, it is often possible to exchange views first, then take in new evidence, and then exchange views again. This alternative usually combines the advantages of *exchanging views first* and *exchanging views last*, as I will point out in footnote 14.

each other.⁶ Surprisingly, however, this does not imply that they should proceed as in (5)–(7), since their case is special.)

3 The Optimal Order of Evidence Acquisition

The information one gets from becoming aware of a peer disagreement is not simply an additional piece of (first- or higher-order) evidence; rather, it has a *retrospective aspect*.⁷ This means that, unlike other first- or higher-order evidence, the evidence from disagreement calls into question the way we have evaluated our original evidence.

Take, as an example, a history professor who seems to recall that Julius Caesar died in 42 BC. Before mentioning the date in his lecture, however, he checks it in an encyclopedia and reads, to his utmost astonishment, that Caesar died in 44 BC. As the new piece of information is presumably more reliable than his memory evidence, he should revise his original belief. This does not imply, however, that it was irrational for him to believe that Julius Caesar died in 42 BC *before* he looked up the date; on the contrary, we can assume that our history professor is usually quite good at remembering historical dates and thus justified in trusting his memory, as long as there are no defeaters.

We can make an analogous observation for acquiring new pieces of higher-order evidence: If the history professor learns that in the recent past he has been far less reliable at remembering historical dates than he once was, he should stop trusting his original belief that Julius Caesar died in 42 BC. But again, this does not imply that it was irrational for him to have this belief *before* he learned about the recent failures of his memory; on the contrary, his then-undefeated evidence supported the view that his memory was to be trusted.

Things are different with evidence from peer disagreement. This evidence raises the question whether one has evaluated the original evidence correctly. The

⁶ ‘Talking to each other’ (or, more precisely, ‘disclosing ones’ respective views regarding some subject matter to each other’) is just what I mean by ‘exchanging views.’ Under the assumptions that the people involved are peers, that their views are different, and that EW holds, it follows that the people involved should revise their beliefs and adopt some kind of middle position. If EW is interpreted in terms of *splitting the difference* (as I presuppose in this and the following section), exchanging views should result in adopting a credence that equals the relevant arithmetic mean; if EW is interpreted as suggested in Section 4, exchanging views should result in becoming agnostic in the way described in that section.

⁷ See Lasonen-Aarnio (2014, p. 317), as well as Rosenkranz and Schulz (2015, pp. 570–72). The term ‘retrospective aspect’ is due to Lasonen-Aarnio. As I will point out below, however, I do not follow her in using it in an externalist way.

fact that, on the basis of the same first-order evidence, one's epistemic peer comes to a different conclusion indicates, according to conciliatory views, that it has never been rational to evaluate this evidence as one did. If our history professor arrives, after careful examination of all relevant data, at the view that Julius Caesar, had he continued to live for another 10 or 20 years, would have ruined the Roman Empire, while a colleague, after carefully analysing the same data, defends the opposite conclusion, both should learn from their disagreement, according to conciliatory views, that they overestimated the conclusiveness of their data and should not have drawn their original conclusions in the first place. Unlike first-order evidence and ordinary higher-order evidence, evidence from peer disagreement calls into question the reasoning on whose basis one formed one's relevant beliefs *before* one acquired the evidence from peer disagreement. This is the retrospective aspect.

Note that I use the term 'retrospective aspect' in an internalist way.⁸ As the history professor who has recently been unreliable at remembering historical dates had, before learning about these failures, no evidence for distrusting his memory, he was, before learning it, internalistically justified in believing that Julius Caesar died in 42 BC. But he had at no time an externalist justification for this belief because his cognitive capacities had not been working well all along.⁹ Understood externalistically, a retrospective aspect is a common feature of higher-order evidence, insofar as higher-order evidence often suggests (perhaps misleadingly) that an external justification was missing all the time. Understood internalistically, a retrospective aspect is a specific feature of evidence from disagreement, insofar as this evidence tells one that one has misevaluated one's first-order evidence, or has overestimated its conclusiveness, and hence was not even internalistically justified in believing what one originally believed (as in the case of the history professor who thinks that Julius Caesar, had he lived longer, would have ruined the Roman Empire). In this internalist understanding, the retrospective aspect marks a categorical difference between non-disagreement evidence and evidence from disagreement.

This categorical difference explains why Bayesian conditionalisation does not work for the former: the newly gained piece of evidence from disagreement cannot just be added to the old evidence; rather, it calls for a reassessment of it. Bayesians can deal with such reassessments, not by conditionalisation, but by calling into question the agents' prior probabilities (see Rosenkranz and Schulz 2015, Section 7). This means, however, that they have to treat evidence from disagreement quite

⁸ I am grateful to an anonymous referee for pressing me on this point.

⁹ Cf. Lasonen-Aarnio's discussion of Elga's hypoxia case (Lasonen-Aarnio 2014, pp. 315–17; see also Elga 2008 for the hypoxia case).

unlike other kinds of evidence. Moreover, a revision of prior probabilities raises the questions how exactly the revision should be conducted, and whether it is, like conditionalisation, a deterministic process, the alternative being that there are several rationally permissible ways of revising the prior probabilities (see Rosenkranz and Schulz 2015, Section 8).

Remember that h is the proposition that the butler is the murderer, and e the proposition that the gardener saw the butler in the billiard room just before the crime happened. If Richard and Siena form different credences towards h , or towards h given e , they are already making a mistake, because the fact that a peer, while evaluating the same evidence, intuitively favours a credence unlike the one oneself intuitively favours implies that the available evidence is not sufficiently conclusive, and that it is not rational to adopt the credence which one tends to find most plausible. One should rather become agnostic about what the right credence is (see Section 4).

Exchanging views may help them to avoid the mistake of overestimating the conclusiveness of their evidence. For if, say, Richard has a specific credence towards h and then notices that Siena has a different credence towards h , he learns that it has never been rational for him to evaluate the evidence as he did, and that he should rather have been agnostic.

But *when* should Richard and Siena exchange views? If they exchange views before they interview the gardener and come to know whether e , the information they get from a potential disagreement is that at least one of them has misevaluated the original evidence concerning h (not including e). If, on the other hand, they first learn that e and then discuss their views, the information they get from a potential disagreement is that at least one of them has misevaluated the total evidence concerning h (including e). Hence, they do not get exactly the same information from a potential disagreement. This explains the different result that we get from different orders of evidence acquisition.

What is more, this shows that the information we get from *exchanging views last* is more encompassing: it entails not only how an epistemic peer interprets the original evidence without the new piece that we are about to get, but also how the peer assesses this new piece and how she incorporates it into her view.

Does that mean that Richard and Siena should proceed as in (5)–(7)? Not quite. What complicates matters here is that in (8) and (9), Richard and Siena do not discuss and revise their respective beliefs concerning h , the hypothesis they really care about, but those concerning e and $h \wedge e$. In other words, they exchange views about how likely they take it to be that e is true and that $h \wedge e$ is true. From this, one can easily calculate how likely they should take it to be that h is true given e . If they then come to know that e is indeed true, it is clear what their credence should be

(that is why we could ascribe them the same credence in (10), namely $C_{R/S}^{+D+e}(h) = \frac{3}{5}$). In order to see how to handle this case, let us first look at

Example 2. Let p and q be two probabilistically independent propositions and $C_R(p) = C_R(q) = \frac{3}{4}$, $C_S(p) = C_S(q) = \frac{1}{4}$. Which credence should Richard and Siena then form towards $p \wedge q$ if they are able to exchange views? (Cf. Staffel 2015; Weatherson 2019, pp. 216–17.)

If they first calculate their resulting credences and then exchange views, we have $C_R(p \wedge q) = \frac{9}{16}$ and $C_S(p \wedge q) = \frac{1}{16}$ and finally

$$C_{R/S}^{+D}(p \wedge q) = \frac{1}{2}C_R(p \wedge q) + \frac{1}{2}C_S(p \wedge q) = \frac{5}{16}. \quad (11)$$

If they first exchange views and then calculate, we have $C_{R/S}^{+D}(p) = \frac{1}{2}$ and $C_{R/S}^{+D}(q) = \frac{1}{2}$ and thus $C_{R/S}^{+D}(p \wedge q) = \frac{1}{4}$.

The problem is that it is certainly wrong for Richard and Siena to have credences $\frac{1}{2}$ towards both p and q but credence $\frac{5}{16}$ towards $p \wedge q$, because p and q are assumed to be probabilistically independent, so that a rationally coherent person's credences should satisfy $C(p \wedge q) = C(p)C(q)$. This is only guaranteed if the calculation takes place *after* the exchange of the views.

Arguably, results of calculations or inferences do not qualify as new evidence. But even if they do, they are certainly unlike other pieces of evidence insofar as they are entailed by the original evidence. As a consequence, we only get more information from exchanging views after doing the calculations or inferences if there is disagreement about how to carry them out. We can, however, plausibly assume that Richard and Siena are not split over such things. Therefore, my argument in favour of *exchanging views last* does not apply to EXAMPLE 2.

Quite to the contrary: EXAMPLE 2 shows that Richard and Siena should first exchange views about p and q and then calculate their credences for $p \wedge q$.¹⁰ Otherwise, they would violate $C(p \wedge q) = C(p)C(q)$, and at least when it is obvious that p and q are independent (as it is, for instance, if p is the proposition that Brazil will win the next football world cup, and q the proposition that New Zealand will be the next rugby world champion) this can only be done at the expense of becoming rationally incoherent.¹¹

¹⁰ On a more general level, Wagner gives similar reasons for the thesis that individuals should apply averaging procedures to arrive at consensual probabilities *before* doing a particular calculation, and not afterwards (Wagner 1985, pp. 99–104).

¹¹ *Exchanging views first* is also the strategy to choose if there is an easy way of checking whether an intermediate result is correct, because then the peers can, if they discover a disagreement,

In general, if we apply conciliatory views to results of calculations or inferences, and not only to initial values or premises, we take the risk of increasing whatever mistakes have so far been made in forming beliefs that are in accordance with the available evidence. The reason is that we may add some internal incoherence to the initial errors: If the parties to a disagreement draw inferences from their original controversial credences, they usually commit consequential errors of different extents. Splitting the difference after drawing such inferences then leads, for each party, to an incongruity between their revised credences in the premises and their revised credences in the conclusion. This is also what goes wrong in (11), according to *splitting the difference*: Richard's and Siena's revised credences in p and q do not match their revised credences in $p \wedge q$. In order to avoid this, and to stay rationally coherent, we should, in cases like EXAMPLE 2, exchange views *before* drawing inferences.¹²

Let us go back to EXAMPLE 1 and to the question of whether Richard and Siena should proceed as in (5)–(7). Given that Richard and Siena agree on how the conditionalisation in (10) has to be carried out, they do not gain anything from exchanging views *after* learning that e and updating their beliefs by conditionalisation. If, on the other hand, they exchange views first, they avoid the risk of amplifying potential errors to different degrees and thereby becoming internally incoherent. Hence, if we stay within the limits of *splitting the difference*, Richard and Siena should revise their beliefs as suggested by (8)–(10).

This holds under the assumption that they discuss their beliefs concerning e and $h \wedge e$. More realistically, however, they would discuss, first, their beliefs concerning h , since they would like to know to what degree the other thinks that the butler should be treated as murderer, suspect, or innocent, and second, their beliefs concerning h given e , since they would like to know the other's opinion concerning how strongly e supports h . Then my original argument applies and they should exchange views *after* learning that e . To see the details, consider

Example 3. Assume that Richard thinks that the Butler probably is the murderer, and that the probability would be even higher if he was seen in the billiard room at

check who is right and fix the error before going on. As it is helpful to detect errors as early in the reasoning process as possible, *exchanging views first* is in this case superior to *exchanging views last*.

¹² Note that Christensen's famous restaurant case (see Christensen 2007, p. 193), which he uses to argue for conciliatory views, is unlike EXAMPLE 2 in that the diners start their calculations with the same numbers (the total amount on the bill, the percentage of tip), so that there is only one new piece of information involved, namely that their calculation results differ. Therefore, no question about the order of evidence acquisition arises. In EXAMPLE 2, on the other hand, Richard and Siena start with different credences before they calculate, so that exchanging views would be non-trivial both before and after doing the maths.

an inopportune moment, while Siena thinks it unlikely that the Butler is the murderer and is of the opinion that it does not alter anything whether or not he has been seen in the Billiard room at whatever time. In numbers: $C_R(h) = \frac{3}{5}$, $C_R(h | e) = \frac{4}{5}$, $C_S(h) = \frac{2}{5}$, and $C_S(h | e) = \frac{2}{5}$ (these numbers are not compatible with the ones assumed earlier). Then *exchanging views last* gives us (remember that, because of (4), $C_{R/S}^{+e}(h) = C_{R/S}(h | e)$):

$$C_{R/S}^{+e+D}(h) = \frac{1}{2} C_R^{+e}(h) + \frac{1}{2} C_S^{+e}(h) = \frac{2}{5} + \frac{1}{5} = \frac{3}{5}. \quad (12)$$

If they first exchange views, however, they both form, according to *splitting the difference*, the following credence:

$$C_{R/S}^{+D}(h) = \frac{1}{2} C_R(h) + \frac{1}{2} C_S(h) = \frac{3}{10} + \frac{2}{10} = \frac{1}{2}. \quad (13)$$

This would obviously change their credences towards h given e . For example, Siena, who thinks that e does not have any effect on h , should now revise her original credence $C_S(h | e) = \frac{2}{5}$ to $C_S^{+D}(h | e) = \frac{1}{2}$. And Richard, who thinks that e does have a significant positive effect on h , should nevertheless lower his original credence $C_R(h | e) = \frac{4}{5}$ a bit, since he now takes h to be slightly less likely; hence, $\frac{1}{2} < C_R^{+D}(h | e) < \frac{4}{5}$.

Again, we get different results for *exchanging views first* and *exchanging views last*: $C_{R/S}^{+e+D}(h) \neq C_{R/S}^{+D}(h | e)$. (Moreover, it is unsure whether $C_{R/S}^{+e+D}(h) \neq C_R^{+D}(h | e)$, because we do not know the exact value of $C_R^{+D}(h | e)$). In addition, $C_S^{+D}(h | e) \neq C_R^{+D}(h | e)$. This time, my argument for *exchanging views last* applies: (12) gives us the correct result because it enables Richard and Siena to consider as well what the other thinks about e .¹³

To sum up: If (i) epistemic peers get two additional pieces of evidence for or against some hypothesis h , (ii) one of those pieces consists in learning that they have different credences concerning h , while the other is just a normal piece of

13 Wilson points out that we might get e and D simultaneously (e.g., if peers disagree about whether they disagree about more than one proposition). Then there is no given order of evidence acquisition, and it looks ‘like there can be no good reason to prefer one order over another’ (Wilson 2010, p. 325). Two things should be said in reply: First, the potential fact that there is no order of evidence acquisition does not mean that we cannot decide in which order we should process the two pieces of evidence. If we have independent reason to prefer a specific order (as I think we have), nothing hinders us from processing the different pieces of evidence in this order. Second, in cases of simultaneous evidence acquisition, the information the peers get from the disagreement is that at least one of them has misevaluated the original evidence, *not including e*. So this is one of the unfortunate cases in which we do not get the optimal information. Hence, we should incorporate e and then try to exchange views again.

evidence, and (iii) they are able to choose the order by which they take in the evidence, then the epistemic peers are well advised to exchange views *last*, that is, to consider the normal piece of evidence first and the evidence from disagreement afterwards. If (i) and (iii) hold and if (ii') one of the two additional pieces of evidence consists in learning that they have different credences concerning h , while the other consists in an uncontroversial calculation or inference, then the epistemic peers are well advised to exchange views *first*, that is, to first consider the evidence from disagreement and then do the calculation or inference.¹⁴

Throughout this and the foregoing section, we presupposed that *splitting the difference* is the correct interpretation of conciliatory views. In the next section, we question this presupposition.

4 Disagreement-Induced Suspension of Belief and Imprecise Probabilities

As pointed out at the beginning of Section 3, the information we get from peer disagreement is that at least one peer has misinterpreted the evidence. This does not mean that all credences that deviate from the mean value are wrong; it only means that we do not know which, if any, of the credences is correct. Hence, a more natural interpretation of EW than *splitting the difference* is what I will call *spreading the difference*: we should become agnostic about which credence is correct, rather than adopt a specific credence that is supposed to reflect our agnosticism.

In a bit more detail, *spreading the difference* requires the epistemic peers to entirely withhold judgment regarding propositions that are controversial between them; all of the peers' original credences – and, for plausibility's sake, all intermediate values as well (but see note 16 below) – should be considered equally good options, between which one should not choose. Hence, the kind of suspension of belief that *spreading the difference* suggests is of higher-order insofar as it does not identify suspension of belief with a specific point or area in the credence interval $[0, 1]$. There may also be a kind of suspension of belief that consists in having a

¹⁴ *Exchanging views first, then taking in new evidence, and then exchanging views again* combines the advantages of both strategies: the epistemic peers are neither in danger of exacerbating past mistakes, nor do they fail to take notice of the peer's evaluation of the new piece of evidence. The only drawback of that strategy is that the effort of exchanging views twice is often unnecessary. For instance, in EXAMPLE 2 it is uncontroversial among Richard and Siena how to do the calculation, so that there is no point in exchanging views after that, whereas in EXAMPLE 3, there is no point in exchanging views before taking in new evidence, since a final exchange of views is sufficient to learn how one's peer assesses all relevant pieces of evidence.

credence of $\frac{1}{2}$, or a credence between, say, $\frac{1}{3}$ and $\frac{2}{3}$; but this is not the kind required by *spreading the difference*.

Take our question from the beginning, whether Pete will drink more than three beers tonight. Richard's credence here is $\frac{1}{4}$ and Siena's $\frac{3}{4}$. According to *spreading the difference*, Richard and Siena learn from a disclosure of their disagreement that the reasonable range of credences extends at least from $\frac{1}{4}$ to $\frac{3}{4}$ – and this is all they should believe about Pete's drinking behaviour after discussing the matter. Put differently, *spreading the difference* comes down to the following: the specific credence that the epistemic peers' first-order evidence suggests lies somewhere in the closed interval that is spanned by the peers' distinct credences, but due to our higher-order evidence from disagreement, we have no clue where in this interval it lies.¹⁵

We can model *spreading the difference* by using *imprecise probabilities* (cf. Elkin and Wheeler 2018). Instead of representing an epistemic subject's belief state by a single probability function C , which maps beliefs to real numbers, we can represent it by a (non-empty) set of probability functions \mathbf{C} , the so-called *credal set*. In the case of peer disagreement, we use credal sets to model the higher-order kind of suspension of belief that results from *spreading the difference*; whether or not the peers' original credences are given by ordinary probability functions or credal sets does not matter. For instance, after discussing the matter Richard and Siena should not form a specific credence about Pete's thirst for beer tonight; as the totality of their evidence is inconclusive, they should be agnostic, and their correct belief states are then best represented by the closed interval $\left[\frac{1}{4}, \frac{3}{4}\right]$.¹⁶

¹⁵ What I call *spreading the difference* is similar to an interpretation of EW suggested in Elkin and Wheeler (2018). The difference is that Elkin and Wheeler allow for adopting any of the peers' credences, while I take the retrospective aspect of evidence from disagreement to indicate that one should abstain from adopting any specific credence. In more detail: If I am right in thinking that the fact that one's epistemic peer comes to a different conclusion, given the same evidence, indicates, according to conciliatory views, that it was never rational to evaluate this evidence as one did, then *spreading the difference* cannot allow for adopting any specific credence within some interval. After all, if it was never rational to adopt any such credence, why should it become rational just when a disagreement arises? (Note that Elkin and Wheeler's justification for their interpretation of EW is different from mine. This explains why they do not discuss the tension between their interpretation and the retrospective aspect.)

¹⁶ Often, but not always, the following holds: when a rational person has credence x towards some proposition p , and another rational person has credence y towards p , then all credences in the closed interval $[\min(x, y), \max(x, y)]$ are credences that rational people might have towards p . This does not hold in cases where there is an easy way of checking who is right, for example by calculating again. Here, the first-order evidence clearly supports a specific result, but the epistemic peers do not know which, because of the higher-order evidence, which shows that at least one of

Let us apply this model to the examples that we discussed in Sections 2 and 3. In those sections, my aim was to present counterexamples; for this purpose, it was helpful to use concrete numbers. Now my aim is to disclose deeper interrelations; it is thus best to use variables instead of numbers to increase generality. Let us start with the observation that (5)–(7) lead to a different result than (8)–(10).

Example 1*. Assume that $C_R(e) = a$, $C_S(e) = b$, $C_R(h \wedge e) = c$, and $C_S(h \wedge e) = d$. Hence, since $C_R^{+e}(h) = \frac{C_R(h \wedge e)}{C_R(e)} = \frac{c}{a}$ and $C_S^{+e}(h) = \frac{C_S(h \wedge e)}{C_S(e)} = \frac{d}{b}$, *spreading the difference* yields the following credal set for *exchanging views last*:

$$\mathbf{C}_{R/S}^{+e+D}(h) = \left\{ C \mid C \in \left[\min\left(\frac{c}{a}, \frac{d}{b}\right), \max\left(\frac{c}{a}, \frac{d}{b}\right) \right] \right\}. \quad (14)$$

For *exchanging views first*, we get

$$\mathbf{C}_{R/S}^{+D}(e) = \{C \mid C \in [\min(a, b), \max(a, b)]\}, \quad (15)$$

$$\mathbf{C}_{R/S}^{+D}(h \wedge e) = \{C \mid C \in [\min(c, d), \max(c, d)]\}, \quad (16)$$

and from this¹⁷

$$\mathbf{C}_{R/S}^{+D}(h \mid e) = \{C_1/C_2 \mid \min(c, d) \leq C_1 \leq \max(c, d), \quad (17)$$

$$\min(a, b) \leq C_2 \leq \max(a, b), C_2 > C_1\}$$

$$= \left\{ C \mid C \in \left[\min\left(\frac{c}{a}, \frac{c}{b}, \frac{d}{a}, \frac{d}{b}\right), x \right], \text{ where } x =$$

$$\max\left(\frac{c}{a}, \frac{c}{b}, \frac{d}{a}, \frac{d}{b}\right) \text{ if this is } < 1 \text{ and } 1 \text{ otherwise} \right\}.$$

In the last expression, ‘1’ is included because it might happen that, for example, $b < c$, while credences must of course not be greater than 1.

them got it wrong. The totality of the evidence (first-order and higher-order) then suggests for each of the peer’s different results that we should be agnostic about whether the first-order evidence supports credence 0 or credence 1. In other words, before checking, our credence should be given by the pair $\{0, 1\}$, not the credal interval $[0, 1]$ (This will not lead to significant changes in the examples below. Mostly, one just needs to replace square brackets by curly brackets. Only in (21), ‘ $[\min(ab, ad, cb, cd), \max(ab, ad, cb, cd)]$ ’ must be replaced by ‘ $\{ab, ad, cb, cd\}$ ’; as a result, (20) and (21) always differ.)

¹⁷ Standardly, one’s credence in h after learning e is given by the set of conditional probabilities, e.g., $\mathbf{C}(h \mid e) = \{C(h \mid e) \mid C \in \mathbf{C}\}$. In (17), I use this standard assumption as well as the definition of conditional probability.

It follows that $\mathbf{C}_{R/S}^{+e+D}(h) = \mathbf{C}_{R/S}^{+D}(h | e)$ if and only if

$$\min\left(\frac{c}{a}, \frac{d}{b}\right) \leq \min\left(\frac{c}{b}, \frac{d}{a}\right) \quad (18)$$

$$\text{and} \quad \max\left(\frac{c}{a}, \frac{d}{b}\right) \geq \max\left(\frac{c}{b}, \frac{d}{a}\right). \quad (19)$$

Since we can without loss of generality assume that $a < b$, the conjunction of (18) and (19) is violated iff $d > c$.¹⁸ In other words, if $d \leq c$, as in EXAMPLE 1, either order of evidence acquisition yields the same result (according to *spreading the difference*).

As argued above, Richard and Siena should first exchange views. So (14), not (17), states the correct credal set. The reason is that the calculation carried out in (15)–(17), which is by itself uncontroversial between Richard and Siena, may increase errors inherent in the initial values. EXAMPLE 1* shows that this happens iff $d > c$ (given that $a < b$). In other words, if the credences of one of the peers (in our case, those of Siena) towards both e and $h \wedge e$ are higher than the respective credences of the other peer (recall that b and d are Siena's credences towards e and $h \wedge e$, respectively, and a and c Richard's), the credal sets stated in (14) and (17) differ, and this is because errors get amplified if the peers do not first exchange their views as in (14), but rather proceed as in (15)–(17).

Example 2*. Let p and q be two probabilistically independent propositions and $C_R(p) = a$, $C_R(q) = b$, $C_S(p) = c$, and $C_S(q) = d$. If Richard and Siena first calculate their credences for $p \wedge q$ and then exchange views, they arrive at $C_R(p \wedge q) = ab$ and $C_S(p \wedge q) = cd$, respectively. *Spreading the difference* then yields

$$\mathbf{C}_{R/S}^{+D}(p \wedge q) = \{C | C \in [\min(ab, cd), \max(ab, cd)]\}. \quad (20)$$

If they first exchange views and then calculate, we have $\mathbf{C}_{R/S}^{+D}(p) = \{C | C \in [\min(a, c), \max(a, c)]\}$ and $\mathbf{C}_{R/S}^{+D}(q) = \{C | C \in [\min(b, d), \max(b, d)]\}$, so that

18 Proof: If $0 < a < b \leq 1$, then $\frac{c}{a} > \frac{c}{b}$ and $\frac{d}{a} > \frac{d}{b}$. With regard to (18), consider first the case that $\frac{d}{b} < \frac{c}{a}$. Then (18) is satisfied because $\frac{d}{b} \leq \frac{d}{a}$. If, on the other hand, $\frac{c}{a} \leq \frac{d}{b}$, then $\frac{c}{b} < \frac{c}{a} \leq \frac{d}{b} < \frac{d}{a}$, and (18) is violated. However, given $a < b$, $\frac{c}{a} \leq \frac{d}{b}$ can only hold if $d > c$. With regard to (19), consider first the case that $\frac{c}{b} \geq \frac{d}{a}$. Then $\frac{c}{a} \geq \max\left(\frac{c}{b}, \frac{d}{a}\right)$, and (19) is satisfied. If, on the other hand, $\frac{c}{b} < \frac{d}{a}$, then (19) is violated iff $\frac{d}{a} > \frac{c}{a}$, that is, iff $d > c$. In sum, given $a < b$, the conjunction of (18) and (19) is violated iff $d > c$ (If $a = b$, it is easy to see that both (18) and (19) hold.)

$$\mathbf{C}_{R/S}^{+D}(p \wedge q) = \{C | C \in [\min(ab, ad, cb, cd), \max(ab, ad, cb, cd)]\}. \quad (21)$$

It follows that (20) and (21) agree if and only if

$$\min(ab, cd) \leq \min(ad, cb) \quad (22)$$

$$\text{and } \max(ab, cd) \geq \max(ad, cb). \quad (23)$$

Since we can without loss of generality assume that $a < c$, the conjunction of (22) and (23) is violated iff $b > d$.¹⁹ In other words, if $d \geq b$, either order of evidence acquisition yields the same result (according to *spreading the difference*). This is also illustrated by EXAMPLE 2 (but note that in this example, $a > c$, so that we get the same results for either order iff $b \geq d$).

Again, Richard and Siena should first exchange views. So (21), not (20), states the correct credal set, because the (by itself uncontroversial) calculation that leads to (20) may falsely ignore possibilities that would not have been overlooked if Richard and Siena had disclosed their credences to each other in advance. EXAMPLE 2* shows that this kind of ignorance happens iff $b > d$ (given that $a < c$). This means that possibilities get overlooked if the credence of one of the peers towards one proposition (in our case, Richard's credence towards p) is lower than the other peer's credence towards the same proposition, but the first peer's credence towards the other proposition is higher than the second peer's respective credence.

Example 3*. To complete our overview of the different cases, assume $C_R(h) = a$, $C_R(h|e) = b$, $C_S(h) = c$, and $C_S(h|e) = d$. Then we get by *exchanging views last*:

$$\mathbf{C}_{R/S}^{+e+D}(h) = \mathbf{C}_{R/S}^{+D}(h|e) = \{C | C \in [\min(b, d), \max(b, d)]\}. \quad (24)$$

If Richard and Siena exchange views first, they both form, according to *spreading the difference*, the following credence:

$$\mathbf{C}_{R/S}^{+D}(h) = \{C | C \in [\min(a, c), \max(a, c)]\}. \quad (25)$$

As a consequence of (25), they would probably have to change their credences towards h given e in order to keep them plausible. Since we cannot state in general

19 Proof: If $0 < a < c \leq 1$, then $ab < cb$ and $ad < cd$. With regard to (22), consider first the case that $ab \leq cd$. Then (22) is violated iff $ab > ad$, that is, iff $b > d$. If, on the other hand, $cd < ab$, then (22) is violated because $ad < cd < ab$. However, given $a < c$, $cd < ab$ can only hold if $b > d$. With regard to (23), consider first the case that $ab \geq cd$. Then (23) is violated because $cd \leq ab < cb$. However, given $a < c$, $ab \geq cd$ can only hold if $b > d$. If, on the other hand, $ab < cd$, then (23) is violated iff $cb > cd$, that is, iff $b > d$. In sum, given $a < c$, the conjunction of (22) and (23) is violated iff $b > d$ (If $a = c$, it is easy to see that both (22) and (23) hold.)

how these credences will be modified, there are no general results for *exchanging views first*, and thus no general results about the conditions that a , b , c , and d must satisfy in order to make it irrelevant whether e or D is first taken in.

As argued in Section 3, Richard and Siena should exchange their views after each of them has considered the significance of the new piece of evidence e ; they should proceed as in (24). Only then do they also learn how the other one assesses this significance. Whether we split or spread the difference does not affect this reasoning.

We need not presuppose credence talk to make *spreading the difference* work. If belief, disbelief, and (first-order) suspension of belief are the only doxastic attitudes, epistemic peers should take their disagreements as evidence that they should not adopt a specific doxastic attitude, but rather remain undecided between those doxastic attitudes that have originally been held by one of them. By identifying belief with 1, disbelief with 0, and (first-order) suspension of belief with $\frac{1}{2}$, and by using sets of these three numbers instead of closed intervals, we can generate versions of the examples in this section that do not invoke credences. In a similar fashion, we can generate credence-free versions of the examples in the previous sections, although we have to make an additional stipulation and declare whether we round $\frac{1}{4}$ to $\frac{1}{2}$ or to 0, and whether we round $\frac{3}{4}$ to $\frac{1}{2}$ or to 1.²⁰

5 Conclusion

The retrospective aspect of evidence from disagreement explains why it matters whether epistemic peers exchange views first or last, before or after taking in a piece of non-disagreement evidence. It thereby also explains why we cannot use Bayesian conditionalisation for modelling the acquisition of evidence from disagreement, since Bayesian conditionalisation presumes that the order of evidence acquisition is irrelevant.

Moreover, the retrospective aspect indicates two further points. First, epistemic peers should usually discuss their views *after* incorporating non-disagreement evidence since they otherwise lack the information how the other

20 It is harder to give a *spreading the difference* interpretation not only of EW, but of conciliatory views in general. In order to do so, one should somehow mark some credences in a given peer's credal set to signal that these are considered by this peer as more likely to be right, but I do not see how this could be done in a plausible way. Personally, I do not think that this speaks against *spreading the difference*, since I take EW to be by far the most plausible (since least arbitrary) conciliatory view.

interprets this evidence. Exceptions are cases like examples 1, 1*, 2, and 2*, in which we have to incorporate the results of calculations or inferences rather than original pieces of evidence, and in which it is uncontroversial among the peers (and known to be so) how to do the calculations or inferences. In such cases, epistemic peers avoid the risk of increasing mistakes by exchanging views first.

Second, we should abandon the *splitting the difference* interpretation of conciliatory views in favour of the *spreading the difference* interpretation. Given suitable original credences (as in EXAMPLE 1 and EXAMPLE 2), *spreading the difference* lets the problem of the order of evidence acquisition vanish. However, the problem is not dissolved in general, as EXAMPLE 1* and EXAMPLE 2* show. In cases in which *spreading the difference* yields diverging results for distinct orders of evidence acquisition, epistemic peers should proceed as explained earlier and exchange views last in normal cases and first in cases of uncontroversial calculations or inferences.

References

- Christensen, D. 2007. "Epistemology of Disagreement: The Good News." *Philosophical Review* 116 (2): 187–217.
- Elga, A. 2008. *Lucky to Be Rational*. Unpublished Manuscript. <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.404.9537&rep=rep1&type=pdf> (accessed August 29, 2021).
- Elkine, L., and G. Wheeler. 2018. "Resolving Peer Disagreements Through Imprecise Probabilities." *Noûs* 52 (2): 260–78.
- Jehle, D., and B. Fitelson. 2009. "What is the 'Equal Weight View?'" *Episteme* 6 (3): 280–93.
- Lasonen-Aarnio, M. 2014. "Higher-Order Evidence and the Limits of Defeat." *Philosophy and Phenomenological Research* 88 (2): 314–45.
- Moss, S. 2012. "Scoring Rules and Epistemic Compromise." *Mind* 120 (480): 1053–69.
- Rosenkranz, S., and M. Schulz. 2015. "Peer Disagreement: A Call for the Revision of Prior Probabilities." *Dialectica* 69 (4): 551–86.
- Shogenji, T. 2007. *A Conundrum in Bayesian Epistemology of Disagreement*. Unpublished Manuscript. <http://www.fitelson.org/few/few07/shogenji.pdf> (accessed August 29, 2021).
- Staffel, J. 2015. "Disagreement and Epistemic Utility-Based Compromise." *Journal of Philosophical Logic* 44 (3): 273–86.
- Wagner, C. 1985. "On the Formal Properties of Weighted Averaging as a Method of Aggregation." *Synthese* 62 (1): 97–108.
- Weatherston, B. 2019. *Normative Externalism*. Oxford: OUP.
- Weber, M. A. 2017. "Epistemic Peerhood, Likelihood, and Equal Weight." *Logos & Episteme* 8 (3): 307–44.
- Weber, M. A. 2019. *Meinungsverschiedenheiten*. Frankfurt/Main: Klostermann.
- Wilson, A. 2010. "Disagreement, Equal Weight and Commutativity." *Philosophical Studies* 149 (3): 321–6.