

The Commutativity of Evidence: A Problem for Conciliatory Views of Peer Disagreement

GEORGI GARDINER

Rutgers University

Final version appears *Episteme*

Abstract

Conciliatory views of peer disagreement hold that when an agent encounters peer disagreement she should conciliate by adjusting her doxastic attitude towards that of her peer. In this paper I distinguish different ways conciliation can be understood and argue the way conciliationism is typically understood violates the principle of commutativity of evidence. Commutativity of evidence holds that the order in which evidence is acquired should not influence what it is reasonable to believe based on that evidence. I argue that when an agent encounters more than one peer, and applies the process of conciliation serially, the order she encounters the peers influences the resulting credence. I argue this is a problem for conciliatory views of disagreement, and suggest some responses available to advocates of conciliation.

In this paper I outline a problem for paradigm conciliatory views of peer disagreement. I argue the way conciliatory views are often understood violates commutativity of evidence. First I sketch peer disagreement and explain various versions of conciliationism. I do not weigh in on whether ultimately we should endorse conciliationism; instead I explain how if a person consults multiple peers, and applies conciliationism serially, the resultant doxastic attitude is sensitive to the ordering of consultations. Since ordering does not affect the evidence acquired, conciliation violates commutativity of evidence. I conclude by evaluating some responses available to conciliationism.

Peer Disagreement

Suppose you are confident your local cinema is showing the film *The Hobbit* tonight. You consult your friend, who believes the local cinema is not showing *The Hobbit* tonight. Your friend is equally diligent and possesses equally good memory and judgement. Indeed before encountering the disagreement you judged your friend equally likely to be right about the film schedule as you. She does not seem drunk, tired, or otherwise impaired. You share your evidence; the disagreement remains. What should you do? Given you were confident in your belief before encountering your friend, are you justified in remaining confident?¹

¹ As noted by Bryan Frances (2010: 72), we should distinguish between the questions of what your opinion should be after the disagreement, and what your opinion should be after the disagreement given what it was before the disagreement. This distinction is important because the two questions yield different answers if your prior credence is faulty.

This is the question of peer disagreement and several responses have been proposed.² Before I introduce these responses it is helpful to clarify some aspects of the debate. In this section I sketch peerhood, belief, and disagreement. Firstly, what is epistemic peerhood?³ One proposal holds two people, A and B, are epistemic peers with regard to p, iff A and B have considered equal levels of evidence that bear on p, have comparable skills, virtues, impairments, and training with regard to whether p. A stronger version of this proposal holds A and B are peers only if they have considered identical evidence that bears on p. An alternative approach holds epistemic peers with regard to p are those who are equally likely to be correct about whether p.⁴ According to this proposal, for instance, a mathematical savant might be peers with a typical undergraduate augmented with a calculator – although they do not share evidence, skill, or training, they are equally likely to be correct. The essence of peerhood is that there is some symmetry in epistemic position, such that favouring one view seems unwarranted. In this essay I employ the “likeliness to be right” conception of peerhood; my arguments apply to other conceptions.

Secondly, what is belief? According to the tripartite conception of belief, belief is non-graded: S either believes p, disbelieves p, or suspends judgement. An alternative conception holds belief comes in degrees of credence.⁵ The issue of peer disagreement arises for both conceptions of belief and I show the non-commutativity result obtains for proposals that employ either conception.

Thirdly, what is disagreement? One gloss holds disagreement occurs iff agents possess beliefs with inconsistent contents. Since it does not employ credences, this conception comports with the tripartite conception of belief. But arguably other kinds of disagreement can arise: perhaps A believes p with great confidence whilst B believes p with less confidence. Similarly, A might believe the evidence strongly supports p whilst B believes the evidence only weakly supports p. MacFarlane glosses disagreement as when A and B have non-cotenable attitudes.⁶ MacFarlane’s gloss provides a broad catch-all, since it covers divergences in both attitude and content. Note that non-cotenable attitudes might be too permissive a conception of disagreement, however, since it admits cases that arguably ought not qualify as disagreement. If A possesses credence 0.8 that the cinema is showing *The Hobbit* tonight whilst B possesses credence 0.9, their attitudes are not co-tenable, but since they do not disagree about whether the film is showing, this is not a paradigm case of disagreement. I do not commit to any particular conception of disagreement; instead I explain the non-commutativity objection with uncontroversial cases of disagreement. However we ultimately characterise disagreement, the objection applies.

² See, for instance, Christensen (2007, 2009, 2010); Elga (2007, 2010); Feldman (2006, 2007); Kelly (2006, 2008).

³ A further question about peerhood is whether the relevant notion of peerhood concerns whether A and B are *in fact* peers or whether A *judges* B to be her peer. For ease of exposition I use the latter conception; my argument applies mutatis mutandis to the former conception.

⁴ Kelly (2006) understands peerhood as having comparable virtues. Kelly (2010) understands peerhood using track records. Elga (2007) understands peerhood using likelihood of being correct.

⁵ A third conception holds that we should understand belief primarily as a comparative notion. This view argues that belief facts have the form S believes p more than q. See, for instance, Hawthorne (2009).

⁶ MacFarlane (draft: 6).

The Positions

One position, call it steadfastness, says that on encountering peer disagreement A is (sometimes) justified in maintaining her prior doxastic attitude. Steadfastness holds that peer disagreement does not (always) require a change in attitude. Some support for steadfastness arises from the fact that in many domains of inquiry, such as politics, philosophy, and religion, there is widespread disagreement. If epistemic normativity requires that on encountering peer disagreement A reduce her confidence, the widespread disagreement in these domains means A would lose many of her beliefs; spinelessness would result. Another motivation for steadfastness is that in controversial domains we should expect widespread dissent. Given this expectation, it is implausible that encountering an instance of disagreement provides *extra* reason to revise belief. When an agent forms her belief in a controversial domain, the expected dissent can be taken into account when forming the belief. It is unexplained why meeting any particular dissenter would thereby generate a reason to change her view. But notice that these reasons, whilst potent when considering controversial topics in politics, philosophy, and religion, are not so applicable to other forms of disagreement; we do not typically expect disagreement about mundane practical matters such as the local cinema schedule, and reducing credence in these cases of disagreement will not typically result in spinelessness.

Steadfastness faces objections.⁷ Firstly, given A and B are peers it is equally likely that A made an error as B. But then it seems implausible that A is justified in privileging her own judgement in this case. As Elga argues, it is absurd for A to maintain – in the face of disagreement – that she is right unless she has some additional, independent evidence. But this is precisely what steadfastness requires. Secondly, steadfastness leads to easy knowledge. By repeated peer disagreement with B, and thereby repeated steadfastness, A can justifiably believe that she is more reliable than B. But it should take more than mere disagreement to justify the belief that she is more reliable; it requires some additional evidence that when they disagree, A is typically correct. Thirdly, it seems steadfastness legitimises ignoring evidence. The fact of disagreement with someone antecedently considered a peer is arguably some evidence that the agent was wrong about whether p. The fact that B believes not-p, furthermore, is evidence in support of the proposition not-p itself.⁸ In a normal case of testimony, B's assertion that not-p is evidence which can justify A's believing that not-p. The steadfastness view must explain why when A possesses a prior belief that p, B's assertion that not-p does not transmit similar evidence and does not call for doxastic revision.

An alternative set of views, call them conciliatory views, argue that on encountering peer disagreement, A should adjust her view.⁹ There are several ways for A to alter her view. One view, endorsed by Christensen, holds on encountering disagreement A should suspend judgement about whether p. An alternative view holds A should sometimes alter

⁷ The “spinelessness” motivation, and “easy knowledge” objection are both found in Elga (2007). See also Christensen (2007), Feldman (2006, 2007).

⁸ See Christensen (2007) for argument that the fact of disagreement itself is evidence. See Feldman (2006, 2007) for argument that B's believing not-p is evidence for A that not-p.

⁹ Note I use the term “conciliatory” slightly differently from some uses in the literature. Some theorists use it narrowly, to refer only to views that hold on encountering peer disagreement A should reduce her confidence in the disputed claim (see, for example, Elga 2010). Elga's gloss is too narrow, however, since it excludes those conciliatory views that hold you should sometimes alter the *content* of your belief to align better with your peer's, and those that hold when your peer's credence is higher than yours you should sometimes *increase* your credence.

the *content* of her belief to align better with the content of her peer's belief. A third view holds A should alter her credence to be more consonant with her peer's.¹⁰

The first conciliatory suggestion has some virtues. The correct action on encountering peer disagreement might be further inquiry, and the suspension view can explain this: we should undertake inquiry because we ought to suspend judgement. But there are weaknesses. Suspension seems a radical response to disagreement – especially in domains where we expect disagreement, such as philosophy, religion, and politics – and widespread spinelessness would result. Suspension faces difficulty countenancing the symmetry of peer disagreement, furthermore, since if A believes *p*, and B withholds judgement, then the suspension view recommends A adjust her attitude, whilst B retains hers. Conciliation thereby does not treat A and B equally, since the suspension approach privileges B's attitude.¹¹

The second suggestion enjoys support from common sense and real-life examples. Suppose I estimate a bear's weight as 400kg, and my flatmate retorts, "It looks like 500kg to me", if we are similarly positioned to judge, a natural response is that I should think I under-guessed and so think the correct answer lies between our guesses. Whilst this means my confidence the bear weighs 400kg will indeed decrease, the change in credence does not best capture my doxastic change. A better way to capture the process is positing a change in content.

The content-adjusting approach has not been discussed explicitly in the literature, perhaps because the procedure has limited applicability: it is only appropriate where the proposition is one where "splitting" is viable. It is unclear how to split the content of "the cinema is showing *The Hobbit*", for instance, and suppose A thinks the film starts at 7pm and B thinks it starts at 8pm, whilst there is a natural way to split the content, by believing the film starts between 7pm and 8pm, this may be unreasonable if A and B also know the film starts on-the-hour. The content-adjusting view holds that *where appropriate* one should alter the content of her belief towards the content of her peer's, and otherwise to alter confidence. Note that the content-adjusting and suspension views, unlike the credence-adjusting view, can be endorsed without positing credences.

The third suggestion – that A should adjust her credence to align better with B's credence – is the most discussed conciliatory approach and yields a spectrum of views.¹² The conciliatory extreme holds A should adjust her credence to halfway between her prior credence and her peer's. Call this the "averaging view".¹³ More steadfast versions of peer disagreement hold you ought adjust credence slightly towards your peer's, but not halfway – you can privilege your own judgement on encountering peer disagreement.¹⁴

¹⁰ Steadfast views and conciliatory views do not exhaust the available options. The right reasons view, for instance, argues that when A encounters peer disagreement, the correct response depends on how successfully her prior credence reflected the evidence. The view denies the apparent symmetry between the epistemic positions of A and B. See Kelly (2006: 180; 2010).

¹¹ For further criticisms of belief-suspension understandings of conciliation see Kelly (2010).

¹² Christensen (2009: 186). Note that the content-adjusting can yield a similar spectrum of views.

¹³ Wilson (2010) argues that averaging conciliationism is the most plausible interpretation of the position advanced in Elga (2007).

¹⁴ See, for instance, the "extra weight view" explored in Elga (2007).

A spectrum of credence-adjustment can also be generated by expanding the conciliatory approach to peer disagreement to encompass non-peer disagreement. Non-peer disagreement includes disagreement with a person you antecedently consider more likely to be right than you, but not sufficiently expert that you wholly defer to her, and disagreement with a person you consider less likely to be right than you, but whose opinion you do not disregard entirely. A natural expansion of the motivations driving credence-adjusting conciliation hold when you disagree with someone who is more likely to be right than you, you alter your credence more than halfway towards her credence, and when you disagree with someone who is less-than-a-peer, but whose opinion is not wholly unreliable, you move your credence towards theirs slightly, but not halfway. Call views that say you should adjust your credence, but not to halfway, “weighted” credence-adjusting.

A Problem for Conciliation

In this section I introduce a problem for credence-averaging conciliationism. In the next section I generalise this result to weighted credence-adjusting and content-adjusting conciliationism. I argue the approaches fail commutativity: when A meets multiple peers, and conciliates accordingly, the order of the encounters influences the final credence recommended by the theory.

The commutativity of evidence principle states that:

To the extent that what it is reasonable for one to believe depends on one’s total evidence, historical facts about the order in which that evidence is acquired make no difference to what it is reasonable for one to believe.¹⁵

The principle holds that the order in which one acquires evidence should not affect the epistemic import of that evidence, provided the same total evidence is gathered. The commutativity of evidence principle is widely held and very plausible. Indeed it is often employed as a benchmark by which to evaluate other views and principles,¹⁶ and denying the principle would be radical. The problem is that conciliationism, at least the versions explored above, seems to violate the commutativity of evidence. To see why consider the following case.

Suppose Anna has credence 0.5 that the cinema is showing *The Hobbit* tonight. Anna consults her peer, Beccy, whose credence is 0.1. The averaging view says Anna’s new credence should be 0.3. Anna updates accordingly. Later Anna sees Claire, who is also Anna’s peer with respect to the cinema’s schedule. Claire’s credence is 0.9. According to the averaging view Anna’s credence should increase to 0.6.

Now consider Anna’s counterpart, Anna'. Like Anna, Anna' has prior credence 0.5 that the cinema is showing *The Hobbit*. She meets Claire'. Like Claire, Claire' has credence 0.9, so according to the averaging view Anna' should adjust her credence to 0.7. Anna' then meets Beccy', whose credence is 0.1. According to averaging, Anna' should alter her credence to 0.4. Table One, below, summarises these numbers:

¹⁵ *Handbook for the History of Logic*, vol. 10: *Inductive Logic* ed. by Dov Gabbay, Stephan Hartmann and John Woods. (2010: 451).

¹⁶ See, for instance, Lange (2000) ‘Is Jeffrey Conditionalization Defective By Virtue of Being Non-Commutative? Remarks on the Sameness of Sensory Experience’ *Synthese* 123, 393-403. I return to the plausibility of the commutativity of evidence principle below.

Table One

Person	Original credence	Credence of 1 st consultant	Credence after 1 st consultation	Credence of 2 nd consultant	Resultant credence
Anna	0.5	0.1	0.3	0.9	0.6
Anna'	0.5	0.9	0.7	0.1	0.4

It seems Anna and Anna' encounter the same evidence: consulting persons with credence 0.9 and 0.1. But the final credence recommended by averaging differs. The result of averaging, when applied serially, is sensitive to the order the evidence is acquired. Since, it is widely assumed, evidence is commutative, this is a problem for the averaging proposal.

The resultant credence is sensitive to ordering because the process of blending credences, in this case through averaging, occurs more than once. Earlier credences are thus blended (submitted to being averaged) more; later credences are blended fewer times and so have a greater influence on the resultant credence. Every credence is averaged, but earlier credences are averaged more times, and so are diluted more. Anna's own prior credence is diluted the most! The problem, in short, is that serial averaging is not commutative. This means the problem will arise in almost all instances where A serially conciliates by averaging the credences.¹⁷

Generalising the Result

It is not only averaging credence-adjusting that generates non-commutativity. Recall that weighted credence-adjusting views recommend adjusting credence on encountering disagreement, but not to the midpoint. Consider a credence-adjusting view that recommends Anna weights her judgement as twice as valuable as her peer's. This means she applies the following formula, where C_1 is her prior credence, C_B is her peer's credence, and C_2 is Anna's resultant credence: $C_2 = ((C_1 \times 2) + C_B) / 3$

If Anna, with prior credence 0.5, meets Beccy, who has credence 0.1, the weighted view recommends Anna's credence change to 0.37. When Anna meets Claire, who has credence 0.9, Anna's credence increases to 0.54. Meanwhile if Anna' meets Claire' first, Anna's credence would move first to 0.63, and then on encountering Beccy' adjusts to 0.46. Weighted credence-adjusting also suffers from non-commutativity. Table Two summarises these figures. Again, this result does not rely on any particular features of the example: serial weighted averaging is not commutative, and so conciliation using weighted averaging violates commutativity.

¹⁷ Updating with averaging-based conciliationism fails to commute with updating using Bayes' Rule. This means if two agents both update on new evidence and conciliate with each other, the order in which they do so affects the resultant credence. This problem has been appreciated for a long time. See, for instance, Genest and Zidek (1986). Jehle and Fitelson (2009) and Wilson (2010) discuss further consequences of this result. The non-commutativity problem articulated in this paper, by contrast, is that averaging-based conciliationism as such violates commutativity. This is a more central problem for conciliationism, since it is located within the view itself, and does not rely on an additional commitment to Bayesian conditionalisation.

Table Two

Person	Original credence	Credence of 1 st consultant	Credence after 1 st consultation	Credence of 2 nd consultant	Resultant credence
Anna	0.5	0.1	0.37	0.9	0.54
Anna'	0.5	0.9	0.67	0.1	0.46

Similarly non-commutativity obtains for content-adjusting conciliatory views. To illustrate, consider an averaging version of content-adjustment, which recommends when a person encounters peer disagreement she, where appropriate, averages the content of her view and her peer's. Suppose Anna judges a bear is 100m away. She consults her peer Beccy. Beccy judges the bear's distance is 80m. Content-averaging recommends Anna "split the difference", and believe the bear is 90m away. Anna consults Claire, who believes the bear's distance is 120m. Anna resultantly believes the bear's distance is 105m.

Meanwhile suppose Anna', believing the bear is 100m away, asks Claire' first rather than Beccy'. Since Claire' thinks the bear is 120m away, content-averaging dictates Anna's belief becomes "the bear's distance is 110m". When Anna' consults Beccy', who judges the bear to be 80m away, Anna's new belief is the bear is 95m away. Again, Table Three summarises these numbers. As with credence-adjusting, serial content-adjusting means the judgement of later consultants has more influence on the resulting credence.

Table Three

Person	Original belief	Belief of 1 st consultant	Belief after 1 st consultation	Belief of 2 nd consultant	Resultant belief
Anna	100m away	80m away	90m away	120m away	105m away
Anna'	100m away	120m away	110m away	80m away	95m away

Responses

In this section I examine some responses available to conciliationism. For ease of exposition I discuss credence-averaging, but my comments apply to other conciliatory views. The first response denies evidence is commutative. This response concedes the evidence is the same, and the only change is ordering of evidence acquisition, but maintains that order can influence which doxastic attitude we should have. Note that by conceding the evidence is the same yet the epistemic import is different, this response denies the uniqueness thesis. The uniqueness thesis states that for a given body of evidence and a given proposition, there is at most one level of confidence that it is rational to have in that proposition given that evidence.¹⁸ Denying the uniqueness thesis by itself is not devastating – it is a controversial thesis – but being inconsistent with the uniqueness thesis may be specifically problematic for conciliatory views, since they are often motivated by it. Furthermore, Kelly (2010) argues some versions of conciliationism

¹⁸ Feldman (2007), Christensen (2009).

entail the uniqueness thesis.¹⁹ If Kelly is correct, this is a problem for the response. We should note, though, that even if denying the uniqueness thesis is not a special problem for conciliationism, denying the commutativity of evidence – especially for such mundane cases of updating – is itself a bold and surprising claim.²⁰ If conciliationism entails the non-commutativity of evidence, the view may lose its putative status as being an intuitive and natural position supported by common sense.²¹

Some cases of serial updating on evidence, furthermore, arguably do fail to commute. Perhaps in special cases the agent receives evidence that influences how she should interpret future evidence, for instance, and so if she were to receive the evidence in reverse order her resultant credence should be different. But even if the commutativity of evidence principle fails in exceptional cases, plausibly it holds for standard cases, including the Anna cases above. It is implausible, in other words, that Anna discovers evidence that influences how she should evaluate future evidence. So even if commutativity of evidence does not hold universally, it seems commutativity should apply in cases like these.²²

Note further that even if advocates of conciliationism deny the commutativity of evidence principle, the cases are still problematic for averaging-based conciliationism. The view recommends a significant difference in Anna's final credences, but the cases do not exhibit features – such as a significant difference in evidence – that warrant such divergent credences. Even if the commutativity of evidence principle is specious, the discrepancy in what is recommended for Anna and Anna' should be worrisome for advocates of the averaging views.

A second response denies that Anna's and Anna's evidence is the same, and so the cases are not failures of commutativity. There are two distinct ways to argue that the evidence is not the same. Firstly, the response might argue she has different evidence in virtue of *her* credences being different. To see why, recall the first case. Anna has the following two states: <credence 0.5, meeting someone with credence 0.1> and <credence 0.3, meeting someone with credence 0.9>. This is different from Anna's two states: <credence 0.5, meeting someone with credence 0.9> and <credence 0.7, meeting someone with credence 0.1>. The response holds she has different evidence because credences are evidence. That credences are evidence gains some support from the fact your prior credence influences what you should ultimately believe. If the only thing that influences what you should believe is evidence, then credences must qualify as evidence. There are problems with this line of argument, however: If credences count as evidence, bootstrapping results. Suppose Anna randomly believes the local cinema is showing *The Hobbit*. If Anna's credences qualify as evidence, in virtue of her belief she thereby possesses evidence that the cinema is showing *The Hobbit*. Since evidence justifies, she may end up with a justified belief. A puzzle for this response, furthermore, is that cases

¹⁹ For discussion of the relationship between conciliatory views of peer disagreement and the uniqueness thesis see Ballantyne and Coffman (2012), Kelly (2010) and Christensen (2009: 758). Feldman (2007) employs the Uniqueness Thesis to support conciliation.

²⁰ Much more can be said about the virtues and costs of denying commutativity as a general constraint on updating, and I cannot do justice to the topic here. For discussion of commutativity as a requirement for updating on evidence, see Field (1978), Jeffrey (1965, 2002), Lange (2000), Wagner (2002).

²¹ See, for instance, Elga (2007), for claims that conciliationism has the virtue of being intuitive and natural.

²² Thanks to Markus Kohl and an anonymous *Episteme* referee for discussion of this point.

where Anna has lower proximal credences (0.5 and 0.3), her final credence is relatively high (0.6). When she has higher proximal credences (0.5 and 0.7), her final credence is lower (0.4). If credences qualify as evidence these numbers are hard to render intelligible.

Alternatively advocates of conciliation can argue Anna's evidence is different across the cases, not in virtue of Anna possessing different credences in the two cases, but in virtue of experiencing different *gradients* between her credence and her peer's. To illustrate consider the following, slightly modified, case. Anna has prior credence 0.2. Her friends Beccy and Claire have credences 0.8 and 0.6 respectively. Where Anna consults Beccy first, and follows the averaging prescription, her credence changes to 0.5 after meeting Beccy and rises to 0.55 after meeting Claire. Where Anna' consults Claire' first, on the other hand, her credence first adjusts to 0.4, and rises to 0.6 after meeting Beccy'. These results are summarised in Table Four.

Table Four

Person	Original credence	Credence of 1 st consultant	Credence after 1 st consultation	Credence of 2 nd consultant	Resultant credence
Anna	0.2	0.8	0.5	0.6	0.55
Anna'	0.2	0.6	0.4	0.8	0.6

Again ordering influences Anna's resultant credence, but arguably Anna and Anna' do not experience the same evidence. To see why, consider Anna's perspective. In the first case she is doubtful about p – her credence is low – then meets someone with high credence, and so the disagreement gradient is high, then she meets someone in almost complete agreement. So Anna's experience is big disagreement, followed by almost no disagreement. In the second case Anna' meets someone with a middling divergence in judgement (a difference of 0.4), followed by someone else with a middling divergence (another difference of 0.4). The cases are different when we look at the credence gradients experienced by Anna.

Arguably this difference qualifies as different evidence. This is because when someone antecedently considered a peer reveals significantly different judgement, this “circumstance of disagreement” might itself be reason to downgrade her, and not conciliate.²³ Perhaps advocates of conciliation can appeal to this, and argue that since in the above cases Anna experiences different credence gradients between her and those she consults, the cases warrant different final credences.

Whilst this suggestion is interesting, it cannot explain the different outcomes. Firstly, in the original cases the gradient between Anna and her peers are similar and do not seem sufficient to explain the large divergence in resultant credence. In the first case, for

²³ See, e.g. Elga (2007). To illustrate, suppose you are in the tropics and it has rained every day. Your credence it will rain tomorrow is 0.9. If based on the same evidence your friend, who you had taken to be a peer with regard to the weather, has a credence of 0.1 that it will rain tomorrow, a natural (and justified) response is to downgrade her as a peer. Plausibly her credence, and the difference between your judgement and hers in this instance, is itself evidence that she is not as good at weather prediction as you previously believed. Even though antecedent to the disagreement you had considered her a peer, you are justified taking her judgement less seriously when you update your credence on the basis of the disagreement.

instance, Anna meets someone with a credence 0.4 lower than hers and then someone with a credence 0.6 higher. For Anna', those gradients are reversed – she first meets someone with a credence 0.4 higher, followed by someone with a credence 0.6 lower. This is a small difference in credence gradients to explain the large difference in outcomes. We can imagine cases, furthermore, where Anna and peers have little divergence in judgement (where Anna, Beccy and Claire have credence 0.5, 0.4, and 0.6 respectively, for instance) yet non-commutativity occurs. If peer downgrading explains all non-commutativity, peer downgrading occurs frequently. But widespread downgrading erodes the motivations supporting conciliation.

A third response to apparent non-commutativity holds that after Anna's first consultation, she loses peerhood with the other person. According to this response, peerhood is fine-grained and so Anna does not meet two peers in succession. Instead she meets a peer and then, since she thereby gains evidence, meets an inferior. Since the second person is not a peer, the response continues, conciliation is unwarranted and so the cases do not challenge the conciliatory view. In response, firstly there are weaknesses with employing a fine-grained, unstable notion of peerhood: if peerhood requires precisely matched evidence few people qualify as peers. Since we rarely know precisely what evidence another person possesses, furthermore, we would rarely, if ever, identify peers. When theories of peer disagreement employ a conception of peer disagreement that does not capture real world disagreement, it raises the question of the relationship between the theory's prescriptions and the normativity of real-world cases.

If peerhood is fine-grained, furthermore, recommending conciliation on encountering peer disagreement applies to few cases, which suggests we must expand conciliation to cases of non-perfect peer meetings. As noted above, a natural expansion of the conciliatory approach holds Anna should employ weighted credence-adjustment. But, as demonstrated above, the non-commutativity objection arises for weighted credence-adjustment. Thus, it seems understanding peerhood as fine-grained makes little progress against the non-commutativity objection.²⁴

Finally, proponents of conciliation might argue that when we encounter multiple peers, rather than considering the peers diachronically, we should average the judgements of the peers synchronically. Thus, when Anna meets Beccy then Claire, rather than first averaging credences with Beccy and then averaging credences with Claire, she should average the three credences together at once. This procedure avoids violating commutativity, since it treats the peers synchronically, and non-commutativity cannot arise synchronically. Some support for averaging synchronically comes from considering that if the three peers were consulting altogether, and after sharing their evidence and reasons they still disagreed about p , a natural conciliatory approach would average their credences amongst the three of them. And so, the thought continues, if this is the correct approach for Anna when her peers, Beccy and Claire, are together simultaneously, this suggests it is the correct approach when Beccy and Claire are encountered at different

²⁴ Advocates of conciliationism might combine a fine-grained notion of peerhood with the claim that conciliationism instructs Anna to update using Bayes' Rule. Since Bayes' Rule is commutative, this process does not violate commutativity. If conciliationism is simply an application of Bayes' Rule, however, it is unclear what the conciliationist thesis brings to the table. The challenge constitutes a dilemma: either conciliationism does not contribute anything over and above a Bayesian understanding of epistemic normativity, or alternatively it can distinguish itself from Bayesianism, but must carve out a way to do so that does not violate commutativity.

times.²⁵ To evaluate this suggestion we must introduce some terminology: “psychological evidence” concerns what other people believe regarding whether *p*, “first-order evidence” is evidence that bears directly on whether *p* without concerning other people’s judgements.²⁶ According to this terminology, when Anna looks at the bear directly she considers first-order evidence; when she consults Beccy to determine what Beccy thinks, she considers psychological evidence.

The synchronic suggestion departs from conciliationism typically expressed in the literature, since in the literature conciliationism is presented as a procedure to apply each time you encounter a single peer.²⁷ This response instead recommends that whenever you encounter peer disagreement you should recall the peers you have encountered with regard to whether *p*, what their credences were at the time of the encounter, and what your own credence is once you abstract the psychological evidence of those peer encounters. You should then average these credences. When recalling your own credence it is important to subtract the influence of previous peer encounters, because otherwise the procedure double-counts psychological evidence.

One weakness of this view is it is difficult to accomplish. If this view is correct epistemic normativity requires Anna to remember peers she has consulted with regard to *p*, their credence in *p* at the time of meeting them, and what her credence is once she abstracts the influence of all previous peer disagreement. But arguably it gets more difficult. To see why, notice that the same results are obtained if, rather than following this synchronic procedure, we employ a diachronic procedure that weights Anna’s own credence more for each person she has conciliated with in the past. If Anna has conciliated with Beccy then, when she meets Claire, Anna can average her view with Claire’s, but double the weight of her own view, since her view has already been blended with Beccy’s. If Anna has conciliated with Beccy and Claire, then when she meets Dana, Anna’s view can be weighted as three times as much as Dana’s. This procedure generates the same results as the synchronic suggestion, since it is arithmetically equivalent. It is similar to the standard averaging procedure for conciliationism discussed above, but by weighting Anna’s view more as she conciliates more, the procedure avoids the non-commutativity caused by repeated blending.²⁸

This procedure highlights how significant the difference between first-order and psychological evidence is for conciliationism: If Anna has consulted with two other people, her view counts as three times as much as her consultant’s view, if she has

²⁵ Note though that whilst synchronic averaging seems intuitive when Anna meets Beccy and Claire in quick succession, it becomes increasingly less intuitive when the interval between the disagreements spans years.

²⁶ This distinction is also found in Kelly (2010).

²⁷ An exception is Kelly (2010).

²⁸ This discussion makes apparent another consequence for conciliationism: How many occurrences of peer disagreement Anna undergoes can make a significant difference to what she ought to believe. But Anna might meet a peer at a time where her credence is in *agreement* with his, even though she might easily have met him at a different time when she disagreed with him. It would be a strange result for conciliationism if the peer’s view only gets weighted if she met him during a period of disagreement, and is not weighted if they met during a period of agreement. This suggests one ought count meetings of *agreement*, in addition to disagreement, and weight them accordingly. Whilst I expect the conciliationist will agree with this consequence, it bears noting that it has been thus far unappreciated. According to standard conciliationism, occurrences of agreement do not impact what one ought to believe, but for this revised conciliationism, they ought have a significant impact.

consulted with three others, her view counts as four times as much. But this raises the question: what if Anna's consultant, Claire, has herself consulted with other people. If Claire has also conciliated with another peer, then Claire's judgement should not be outweighed by Anna's just because Anna has conciliated before. It seems that to be consistent – to consistently recognise the important distinction between first-order and social evidence – the conciliationist ought to say that when Anna's peer has consulted with other people, and so the peer's view itself is a mixture of first-order and psychological evidence, then the peer's view also gets weighted more accordingly. But now we see how the difficulties multiply: to employ this procedure, Anna must not only recall the peers she has consulted with regard to *p*, she must also learn how many peers *her peers* have consulted with, and so to what degree their judgement is influenced by psychological evidence.²⁹ These values seem very difficult to obtain. If this is what conciliationism demands, it is arguably a revisionary view of epistemic normativity and loses its intuitive appeal. It certainly departs from how conciliation is usually understood. And problems multiply: If one's peers consulted the *same* peers then they should not thereby both count more, since this would double-count psychological evidence. But this suggests one needs to know not just how many peer's peers there are, but also their identities. The inputs look increasingly harder to obtain.

But the procedure is not merely difficult to accomplish – many epistemic norms are difficult to accomplish. The problem is that the procedure seems wholly divorced from our actual epistemic practises.³⁰ Our epistemic practises do not seem so concerned with keeping track of numbers of consultations, credences at times, and with averaging. Perhaps our epistemic practises are misguided, but it seems that conciliationism departs further from our epistemic practises than first realised, and so loses its status as intuitive and supported by common sense. And this modified account of conciliationism departs from that articulated in the literature.³¹

Conciliationism has been criticised for allowing psychological evidence (evidence about how other people have weighed the evidence) to drown out first-order evidence (the agent's estimation of what the evidence recommends).³² But this “synchronic” solution generates a much larger version of this problem: Anna's own judgement can be swamped by many other opinions, including the views of people Anna has never encountered. When deciding what to believe, all the encountered peers of Anna's encountered peers will count as much as her own weighing of the first-order evidence; Anna's own estimation of the first-order evidence may contribute just a tiny fraction to her resultant credence. Since peerhood relations can change over time, furthermore, Anna might also blend her own weighing of the evidence with people with whom she has never been peers. Perhaps at some time Anna was peers with Beccy, for instance, who at some other time was peers with Elaine. Elaine might never have been peers with Anna, but

²⁹ Indeed, a regress looms, for presumably the peer's peers themselves may have conciliated with other peers in the past. If they have, then their views ought to be weighted more. But if this is right, by similar reasoning we must consider our peer's peer's peers, and so on.

³⁰ Thanks to Bradford Cokelet for discussion of this point.

³¹ Note too that if our peers' peers first-order judgements are relevant, it is less plausible that the operative notion of peerhood concerns people we *take* to be peers, rather than those who are in fact peers. This is because we may never encounter our peers' peers. The suggestion thus has more merit when we consider actual, rather than perceived, peers. Thanks to an anonymous *Episteme* referee for this insight.

³² Kelly (2007).

according to this version of conciliationism Anna ought nonetheless blend their views. Again, this departs from how conciliationism is usually understood.

Thus this synchronic suggestion seems to lead to a somewhat revisionary version of averaging-conciliationism. It avoids non-commutativity, but at the cost of making epistemic normativity alien to our current epistemic practises. Given the problem of non-commutativity, however, conciliationists should explore this suggestion as a revised understanding of averaging-based conciliationism. Ultimately the viability of conciliationism rests on its being understood in a way that respects commutativity.³³

Bibliography

- Ballantyne, N. and E. J. Coffman. (2012). “Conciliationism and Uniqueness” *Australasian Journal of Philosophy* 90(4): 657–670.
- Christensen, D. (2010). “Higher Order Evidence” *Philosophy and Phenomenological Research* 81: 185–215.
- Christensen, D. (2009). “Disagreement as Evidence: The Epistemology of Controversy” *Philosophy Compass* 4: 756–767.
- Christensen, D. (2007). “Epistemology of Disagreement: the Good News” *Philosophical Review* 116: 187–217.
- Elga, A. (2010). “How to Disagree About How to Disagree” in *Disagreement* (eds.) R. Feldman and T. Warfield, Oxford University Press.
- Elga, A. (2007). “Reflection and Disagreement” *Noûs* 41: 478–502.
- Feldman, R. (2007). “Reasonable Religious Disagreements” in *Philosophers without Gods: Meditations on Atheism and the Secular Life* (ed.) L. Anthony, OUP. 194–214.
- Feldman, R. (2006). “Epistemological Puzzles about Disagreement” in *Epistemology Futures* (ed.) S. Hetherington, Oxford University Press. 216–236.
- Field, H. (1978). “A Note on Jeffrey Conditionalization” *Philosophy of Science* 45: 361–367.
- Frances, B. (2010). “Disagreement” in *Routledge Companion to Epistemology* (eds.) D. Pritchard and S. Bernecker, Routledge. 68–74.
- Genest, C. and J. Zidek. (1986). “Combining Probability Distributions: A Critique and an Annotated Bibliography” *Statistical Science* 1(1): 114–135.
- Hawthorne, J. (2009). “The Lockean Thesis and the Logic of Belief” in *Degrees of Belief* (eds.) F. Huber and C. Schmidt-Petri, Oxford University Press. 49–74.
- Jeffrey, R. (1965). *The Logic of Decision*. University of Chicago Press.
- Jeffrey, R. (2002). *Subjective Probability: The Real Thing*. Cambridge University Press.
- Jehle, D. and B. Fitelson. (2009). “What is the “Equal Weight View?”” *Episteme* 6: 280–293.
- Kelly, T. (2010). “Peer Disagreement and Higher Order Evidence” in *Disagreement* (eds.) R. Feldman and T. Warfield, Oxford University Press. 111–174.
- Kelly, T. (2006). “The Epistemic Significance of Disagreement” in *Oxford Studies in Epistemology* (eds) T. Gendler and J. Hawthorne, OUP. 167–196.
- Lange, M. (2000). “Is Jeffrey Conditionalization Defective By Virtue of Being Non-Commutative? Remarks on the Sameness of Sensory Experience” *Synthese* 123: 393–

³³ Thanks to Branden Fitelson, Jon Garthoff, Alvin Goldman, Daniel Greco, Jeff Ketland, and an anonymous referee for *Episteme*. Thanks also to audiences at Rutgers University, University of Miami, the Princeton Rutgers Penn (PRP) Social Epistemology Workshop and the Bled conference on Evidence, Reliability, and Group Epistemology.

403.

MacFarlane, J. (draft). “Varieties of Disagreement”.

Wagner, C. (2002). “Probability Kinematics and Commutativity” *Philosophy of Science* 69: 266–278.

Wilson, A. (2010). “Disagreement, Equal Weight and Commutativity” *Philosophical Studies* 149: 321–326.