

*Conciliation, Uniqueness and Rational Toxicity*¹

DAVID CHRISTENSEN
Brown University

Abstract

Conciliationism holds that disagreement of apparent epistemic peers often substantially undermines rational confidence in our opinions. Uniqueness principles say that there is at most one maximally rational doxastic response to any given batch of total evidence. The two views are often thought to be tightly connected. This paper distinguishes two ways of motivating conciliationism, and two ways that conciliationism may be undermined by permissive accounts of rationality (those that deny uniqueness). It shows how conciliationism can flourish under certain strongly permissive accounts of rationality. This occurs when the motivation for conciliationism does not come (as is sometimes supposed) from taking disagreement as evidence of one's own rational failings. However, divorcing the motive for conciliating from worries about rationality does not remove a feature of conciliationism that some find troubling: that conciliationism can lead to cases of "rational toxicity," in which the most rational response to one's evidence involves violating some rational ideal.

1. Introduction

Conciliatory accounts of disagreement hold that the disagreement of others often requires one to be considerably less confident in one's beliefs than one would be absent the disagreement; on such views, the levels of confidence that many people seem to have in controversial opinions are, in virtue of the disagreement, not fully rational. Uniqueness principles about rational belief hold that there is at most one maximally rational response to any given batch of total evidence. These two views have struck many people as tightly related.

Some writers have supported uniqueness claims in order to argue for conciliatory views on disagreement.² Their idea, roughly, is that if I have good reason to think that another person is as likely to react rationally to our common evidence as I am, and if there's only one maximally rational response to that evidence, then the person's disagreement is evidence that my initial belief is not the one the evidence supports, which should make me lose confidence in that belief. And in the same vein, but running in the opposite direction, some have argued against uniqueness and taken that to tell against, or undercut the motivations for, certain versions of conciliationism.³ Some have denied that particular conciliatory positions *entail* uniqueness, though they still see some fairly strong connection.⁴

On the other hand, some have argued for conciliatory positions on disagreement without relying on uniqueness.⁵ And others have denied that uniqueness has much to do with conciliatory views of disagreement at all.⁶

In what follows, I will distinguish two sorts of motivations that have been offered for conciliationism, and two different ways in which conciliationist demands may be undermined by permissive accounts of rationality (those that deny uniqueness). In the end, I will side with those who say that the motivations for requiring conciliation (at least some of them) are compatible with strong denials of uniqueness. I'll first map out some of the ways in which certain (though not all) sorts of conciliatory pressure can arise on certain (though not all) strongly permissive accounts of rationality. One upshot of this is that, insofar as one is disturbed (as many seem to be) by conciliationism because it forbids rational confidence in many of one's controversial opinions, one's concerns may not be assuaged by arguments against uniqueness.

I will then turn to examine a question posed by this line of thought, a question which involves a second feature of conciliatory views that can be disturbing. Conciliation sometimes seems to require agents to violate rational rules or ideals, by adopting final opinions which fail in one sense to give their original evidence its due. Consider, for example, the following example:

Logic Problem: I think through a real-world logic problem with my friend, knowing we have long and equally-good track records with this sort of problem. But this time, we disagree about whether a certain logically complex claim *P* follows from certain other claims *A*, *B* and *C*—claims that we both believe very confidently: I (correctly, as it turns out) think it does, and so am highly confident that *P*. She disagrees, and reports rather low confidence in *P*. I conciliate with my friend, and adopt a middling credence in *P*.⁷

Note that in this case, I end up doubting *P* *despite having undisputed evidence—A, B, and C—which in fact entails that P*. Though I correctly appreciated this entailment at the outset, conciliationism required my giving substantial credence to my friend's contrary opinion. Now I think it's plausible that in giving substantial credence to not-*P* while remaining highly confident in *A*, *B* and *C*, I'm violating a rational ideal that logic imposes on my beliefs. Thus it may be argued that, on the conciliatory view, disagreement evidence can turn out to be in a sense *rationaly toxic*: it may put agents in a position where the most rational response to their total evidence violates some rational ideal.⁸ This can seem like an unfortunate result—even a reason for resisting conciliatory views.⁹

Now it turns out that conciliatory demands in *permissive* contexts occur most clearly on accounts where the rational pressure to conciliate does not flow, as it is often thought to, from the disagreement's providing evidence that one's initial belief is irrational.¹⁰ And this may seem to offer an additional possible benefit to the conciliationist. Heeding (possibly misleading) evidence about the rationality of one's beliefs seems like just the sort of thing that would be expected to lead to one's violating rational ideals. So one might hope that in divorcing the motivation for conciliation from an agent's worries about her rationality, one might help avoid this problem. I will argue, though, that the possibly troubling phenomenon persists, even in clearly permissive settings, and even where conciliation is not motivated by an agent's doubts about her rationality.

2. Two Strands of Conciliatory Argument

Nobody thinks that one is required to conciliate in the face of just any disagreement. The literature typically concentrates on people one has (independent of one's views on the disputed issue) good reason to take as *epistemic peers*—as rough equals along certain dimensions of epistemic evaluation.¹¹ One such dimension concerns the evidence the other person has relevant to the disputed issue, and the other concerns how well she forms beliefs on the basis of her evidence. For present purposes, I'd like to put aside the first dimension, and focus on cases where the disagreeing parties have (at least roughly) the same evidence relevant to the dispute. But we should notice that there are a couple of different ways of approaching the second dimension of evaluation—ways which are not always clearly separated. One focuses on the other person's equal likelihood of responding *rationally* to her evidence. On this reading, rational pressure to conciliate stems from having good dispute-independent reason to believe that the disagreeing friend is what might be called a "rationality-peer" on the given issue: one whose opinion is equally likely to be rational.¹² The second way of evaluating the other person's responses to evidence is in terms of her likelihood of responding to that evidence by forming *accurate* beliefs. On this reading, rational pressure to conciliate stems from having good dispute-independent reason to take the the disagreeing friend as what might be called an "accuracy-peer" on the given issue: one whose opinion on the disputed issue one expects to be as likely to be accurate as one's own.¹³

It is perhaps natural to run these two notions of responding well to the evidence together.¹⁴ After all, it's natural to assume that responding rationally to one's evidence and achieving accurate beliefs will go hand in hand. More specifically, when one thinks about a friend who shares all of one's evidence relevant to a certain issue, the two ways of evaluating the friend will come apart only insofar as one thinks that one of you has a way of forming accurate beliefs that goes beyond rationally responding to the evidence (we'll come back to this point below). Nevertheless, if we want to investigate the extent and nature of the connections between conciliationism and uniqueness, it will be useful to separate the two ways of motivating conciliation.

3. Rationality-Based Conciliation and Uniqueness

Now we should note that, even taking the conciliatory pressure to come from rationality-evaluations, there seem to be many accounts of rational belief which deny uniqueness, but allow for a good deal of conciliatory pressure. For example, thinking of belief in a categorical way, one might deny uniqueness by holding that in certain cases evidence can make rational either suspending judgment about P or believing P, but one might still deny that believing P and believing \sim P could both be rational for agents with the same evidence.¹⁵ Since most of the disagreements in, for example, philosophy seem to be of the "believe P/believe \sim P" type, the disagreement of an apparent philosophical rationality-peer would provide one with significant reason to doubt that one's own belief on the disputed matter was rational.

And that might lead to exactly the sort of requirement to withhold belief on controversial philosophical issues that critics of conciliationism often worry about. A parallel point may be made if we think about belief in a degreed way: as long as the range of rationally permissible credences is limited, disagreements involving widely separated credences may provide evidence that one's own reasoning on the disputed matter was rationally defective.¹⁶

An additional point can be made about theories that violate uniqueness because they make rational belief depend on non-evidential agent-specific factors (such as an agent's values, or practical interests). Suppose, for instance, it turned out that rational belief was sensitive to agents' practical stakes in the way suggested by the account of justified belief in Fantl and McGrath (2002). It's far from clear that this view would allow for steadfast responses in typical cases. We might often have excellent reason to believe that the differences of opinion between, say, P-believers and P-withholders in philosophy were not explained by differences in the relevant sort of practical stakes.¹⁷ Supposing that the alternative explanation would involve failure on the part of some philosophers to adopt the doxastic attitude that was rational for them, such disagreements could well often, even typically, give agents significant reason to doubt the rationality of their own beliefs.

So the connection between uniqueness and conciliation is not simple, even when conciliation is motivated by rationality-evaluations of other agents. Nevertheless, for present purposes, I want to concentrate on the strand of conciliationist thought that would seem much less likely to depend on uniqueness—the strand that has agents evaluate the opinions of others in terms of accuracy, rather than rationality.¹⁸

4. Accuracy-Based Conciliation and Uniqueness

As noted above, one of the early defenses of conciliationism (Elga 2007) both based conciliatory pressure on accuracy-based evaluations of the other agent, and disavowed any commitment to uniqueness. And it might seem quite clear that the accuracy-based approach to peerhood could produce conciliatory pressure without relying on claims about the peer's rationality. If one is using one's peer's beliefs as an indicator of worldly states, why think that rationality considerations would be relevant?¹⁹ One does not need a thermometer to be rational in order for its readings to provide reason to revise one's temperature beliefs.

I think this line of thought is basically sound. But it turns out that there are some interesting complexities involved. In this section, we'll see that different sorts of permissive theories of rationality have different implications for the significance of peer disagreement, even when we take peerhood in the accuracy-based sense.

4.1. *An extremely permissive account of rationality*

Insofar as we're trying to separate conciliatory pressure based on considerations of rationality from conciliatory pressure based on considerations of accuracy, we might do well to begin with a toy theory of rationality which puts maximum distance between these two factors. So let's consider the Epistemic Hippie's Theory

of Rationality, often expressed in the admonition, “If it feels good, believe it!” but formulated more precisely as follows:

EHTR: Necessarily, for any subject S , proposition p , and time t , S 's belief that p at t is maximally rational iff it doesn't engender a harsh vibe.

To avoid complications, let us suppose that Abbie believes this theory of rationality, and that the theory is true. We'll suppose that Abbie has a confident belief that P , and encounters disagreement from a friend.

Now it seems clear that Abbie need feel no pressure at all to conciliate with a disagreeing friend who's a peer in the *rationality* sense. On his theory, the disagreement of a rationality-peer doesn't call into question either the rationality or the accuracy of his initial belief.

On the other hand, suppose we consider how Abbie should react to the disagreement of an apparent *accuracy*-peer. Clearly, the *rationality* of Abbie's initial belief is not brought into question by disagreement of a friend he takes to be an accuracy-peer. But since he sees her as an accuracy-peer, her disagreement would seem to call into question the accuracy of his initial belief that P . And this, it might seem, should provide reason for him to lose confidence in that belief. After all, it would seem that Abbie would be concerned with having accurate beliefs, not just with having rational ones. For example, if he learned that what he was just smoking had the effect of giving him wildly inaccurate beliefs about his surroundings, yet ensured that these beliefs engendered no harsh vibe (and hence, were rational), we might well expect him to lose confidence in the truth of those beliefs. It might even seem obvious that he *should* lose confidence in those beliefs.

However, there is a problem with this line of thought—at least if the “should” is taken to be connected to epistemic rationality.²⁰ The problem is that it's not true that Abbie is rationally required to lose any confidence in P . On EHTR, which we're supposing for now to be true, Abbie may maintain his belief with full rationality, so long as doing so produces no harsh vibe. In fact, Abbie seems not to be rationally required to react to any new evidence at all!

Does this mean that Abbie can rationally say something like, “Well, I guess my belief that P is as likely as not to be false. But I don't care, since I'm rational either way”? Not necessarily. There's something (at least close to) incoherent about describing someone as believing that P , but simultaneously believing that his friend's belief that $\sim P$ is just as likely to be true. It may be thought that there's not clearly enough daylight between *believing that P* and *believing that a belief that P is true* for this to make sense. Nevertheless, we don't need to take a position on that issue here. It seems clear that Abbie could simply continue to believe P , and also believe that despite his friend's general accuracy-peerhood, his friend just got unlucky when it came to P . As long as it engendered no harsh vibe, Abbie's belief would remain maximally rational.

This highlights a problem with approaching our question from the perspective of EHTR. EHTR does, as we've seen, undermine conciliationism. And it does this in a way that flows directly from its extreme permissiveness. However, the way in which it undermines conciliationism is extremely crude. On EHTR, an agent

may believe whatever she wants, perhaps simply ignoring the evidence completely, and as long as this engenders no harsh vibe, her beliefs will be fully rational. The account is so flexible that it will be hard to find any phenomena that will put general rational pressure on agents. So while it may provide a dramatic way of separating rationality-worries from accuracy-worries, it does so to a degree that makes optional all of the intuitively rational reactions to information that should raise worries about accuracy of an agent's beliefs.

It turns out, then, that there are two different ways in which permissiveness can undermine conciliationism. The first, illustrated by EHTR, is by being so permissive that it undermines virtually any constraint on rational responses to evidence. But I think it's clear that those who have thought that permissive accounts of rationality were incompatible with conciliationism have had something much less crude in mind. The idea is not that conciliationism is false because anything goes. The trouble is supposed to stem from the thought that, even if there are substantial constraints on rational belief, still, insofar as rationality is permissive, my finding out that someone with my evidence disagrees with me need not impugn the rationality of my initial belief. This suggests that we look at an account that's highly permissive, but somewhat less crazy than EHTR.

4.2. *Subjective Bayesianism*

Let us begin by considering Subjective Bayesianism, which takes rationality to be defined by certain formal properties of the agent's belief-system. A rational agent's beliefs must be probabilistically coherent, and she must change her beliefs only by conditionalization (including, perhaps, Jeffrey's generalization of conditionalization). But those are all the constraints that rationality puts on belief.²¹ Suppose that this is the correct account of rational belief, and suppose that Bruno accepts it. Bruno's view is very permissive: as long as someone stays coherent and changes belief by conditionalization only, Bruno will count her beliefs as rational. Now, lest we run into the problem we had in concentrating on Abbie, let's make Bruno more intuitively sensible in his governing of his own beliefs. Let us suppose that Bruno's own prior probability function is one we would consider generally reasonable. In particular, if he encounters evidence that (as we'd see it) strongly suggests that his beliefs are inaccurate, he'll change his beliefs accordingly.

For example, suppose he's a detective, and becomes highly confident, on the basis of strong evidence, that Jocko committed the crime. Then he gets new evidence indicating that his belief is inaccurate: say, he learns that there was evidence-tampering intended to produce false beliefs in many cases, or that Jocko has an airtight alibi, or that DNA found at the scene matches that of another suspect, who has confessed. So Bruno reduces his confidence in Jocko's guilt. While there is nothing in the abstract framework of Subjective Bayesianism that requires this response of everyone (certain wacky priors would result in such additional evidence *increasing* the agent's confidence in Jocko's guilt), we'll imagine that Bruno's sensible prior builds in the sort of respect for this sort of evidence that mandates loss of confidence in his hypothesis.²²

Now it seems clear that, like Abbie, Bruno may well be (very reasonably) quite untroubled by the disagreement of *rationality* peers. It's part and parcel of the Subjective Bayesian outlook that there are countless maximally rational possible reactions to one's evidence that are nevertheless wildly inaccurate.

In fact, to greater or lesser degrees, it seems that almost any robustly permissive account of rationality will require an agent to think of certain others as having the following characteristics: they have all the same relevant evidence as the agent herself; their beliefs are perfectly rational given that (shared) evidence; and their beliefs are highly inaccurate.²³ There is, one might think, something a bit odd about this combination of attitudes. The oddness comes out in thinking about how the agent with such a view should regard her own beliefs—in particular, how she should answer the question: “What explains my being correct and my friend's being incorrect in the present case?” By hypothesis, it's not that she has some evidence her friend is not privy to. And by hypothesis, it also can't be that she's drawn the more rational conclusion from their common evidence. Moreover, it wouldn't seem reasonable for her to think she got the present issue right by pure luck (after all, we may suppose that she and her friend disagree about many things, and in coherently maintaining these beliefs, she must think that she's right much more often than her friend is). Presumably, what she should think is something like this: that she has some way of forming accurate beliefs *that goes beyond reacting rationally to her evidence*. This sort of view may seem unintuitive.²⁴ But it looks to be something that most permissivists will need to embrace.²⁵

So let's suppose that this is how Bruno sees things: the fact that rationality-peers with his evidence disagree with him gives little cause for concern that his own initial credence was not rational.²⁶ Nevertheless, it seems that Bruno may well be (again quite reasonably) much more impressed by disagreement of *accuracy-peers* (just as he's impressed when he finds out about the evidence-tampering). Let's consider an example:

Bruno has solved many cases correctly, as has Callie. Indeed, they both have long and admirable track-records of solved cases. When looking at the same batch of evidence, they almost always agree. But on the occasions when they've become confident in incompatible conclusions, they've been right (in the sense of being confident of the true conclusion) equally often. They both examine a thick file of evidence, and find themselves confident in conflicting hypotheses. Callie says, “One of us has failed to reach the rationally supported hypothesis. Perhaps, lacking a clue as to which of us it was, we should each reduce confidence in our hypotheses.” But just then, the Epistemology Oracle steps in. “Look, I'm not going to tell you what to believe. But I'll say this: Uniqueness is false, because Subjective Bayesianism is true. And I'll tell you something else: you're both sterling Bayesian agents. You've each almost always had only rational credences, and your initial credences after studying today's evidence file were no exception.”

This still seems to be a case in which Bruno may have ample reason to revise his belief. Bracketing his own reasoning about the present criminal case, he has excellent reason to think Callie equally likely to reach the correct verdict. So if his prior encodes what a conciliationist would take to be the reasonable reaction to

disagreement of this sort, Bruno will take Callie's disagreement as evidence against the accuracy of his belief—her disagreement makes it less likely, by Bruno's own reckoning, that Jocko committed the murder. In fact, Bruno may well be *rationally required* to lose confidence in his initial hypothesis, simply because his prior encodes (what we might think of as) the sensible reactions to this sort of information.²⁷ But Bruno's reason for conciliating with Callie does not stem in any way from Callie's disagreement impugning the rationality of his initial belief. It's seeing her as an accuracy-peer that does all the work, by providing reason for Bruno to suspect that his initial belief, even if it was perfectly rational, may be inaccurate. So in Bruno, we have a case of rationally mandated conciliation in the context of an extremely permissive account of rationality.²⁸

It should be noted, however, that there is an important limitation to the force of this example. While we do have a case where conciliatory rational pressure exists without uniqueness, this case cannot be used to show that *the general position of conciliationism* is tenable without endorsing uniqueness. Conciliationism is supposed to entail that people who hold the sort of confident opinions we often see on controversial matters in areas such as philosophy, politics, economics, etc., should in many cases be much less confident in their opinions than they often seem to be. But the sort of rational conciliatory pressure we saw in Bruno's case clearly depended on specific contingent facts about Bruno: it depended on his particular prior probability function. That certainly doesn't translate into any general mandate for conciliation in the face of (even accuracy-defined) peer disagreement. So the sort of permissiveness embodied in Subjective Bayesianism does seem to be incompatible with the truth of conciliationism, when we see conciliationism as embodying a general rational requirement.

Now it's worth noting that there is a sense in which a *kind* of conciliationism might be true even if Subjective Bayesianism is the correct account of rationality. This sense will depend on the actual distribution of priors. If the vast majority of people have priors like Bruno's, then there may in fact be a widespread rational obligation to conciliate with apparent accuracy-peers. And many people may (again, due to their sensible priors) have good (dispute-independent) reason for believing that those who disagree with them are accuracy-peers. But even if that were true, conciliationism would hold as a contingent fact, not as a matter of general rational principle.

This brings to the fore that Subjective Bayesianism, as a complete theory of rationality, shares important features with EHTR. The constraints that pure Subjective Bayesianism places on rational belief are in a way exceedingly thin; there will be precious few general rational principles at all. This is, after all, a view according to which the evidence that we currently have could rationally support a theory on which Elvis is alive, Barack Obama is an alien, and George W. Bush is Mahatma Gandhi reincarnated.²⁹ It's hardly surprising that such a view of rationality would permit all kinds of intuitively bizarre reactions to peer disagreement. So to see how conciliationism *as a general principle* could be combined with a permissive theory of rationality, we will need to consider an account of rationality that is permissive, but which has more room for substantive general rational requirements.

4.3. *A Sanelly Permissive Account of Rational Credence*

Let us begin by narrowing down pure Subjective Bayesianism, eliminating some intuitively wacky priors. We might eliminate priors that violate the Principal Principle. We might eliminate priors that, given an ordinary person's evidence, endorse the Brain-in-Vat theory over the Ordinary-Physical-World theory. Or priors on which finding green emeralds this year strongly supports finding blue ones next year. And so on. The resulting theory might be called Constrained Subjective Bayesianism.

It's important to notice that there's no reason to think that eliminating intuitively wacky priors would come anywhere close to imposing uniqueness. And of course most critics of uniqueness are not motivated by thinking that our theory of rationality should allow for the belief that Elvis is alive to count as rational on our present evidence. They typically point out that paradigmatically reasonable people seem to reach different conclusions from the same evidence; or they claim that desiderata for hypothesis-choice (such as simplicity, or fit with the data) are subject to different perfectly rational interpretations, and/or different perfectly rational weightings against one another. So Constrained Subjective Bayesianism's narrowing down the range of rational priors to eliminate the wacky ones still allows for a robustly permissive account of rational belief.

Nevertheless, it's also plausible that when we narrow down the range of rationally acceptable priors to the non-wacky ones, we will end up mandating certain sorts of reactions to evidence: ones that help constitute rational believing. For example, all acceptable priors may require that in typical situations where an agent has formed a belief about the perpetrator of a crime, she lose confidence when she comes to learn, say, that the evidence on which she based her belief was planted (or that another suspect confessed and only that suspect's DNA was found at the crime scene, or that she was slipped a powerful judgment-distorting drug just before studying the evidence). And it is also possible that all acceptable priors will require loss of confidence in many situations where the agent learns of the disagreement of others whom the agent has excellent dispute-independent reason to regard as her (equally-informed) accuracy-peers.³⁰ In other words, some ways of constraining Subjective Bayesianism will impose conciliatory demands as a general rational requirement: what applied to Bruno in the previous section would apply to everyone.

So some forms of Constrained Subjective Bayesianism seem to show that at least a fair amount of conciliation with apparent accuracy-peers can be rationally required even in a strongly permissive setting. But before going further, I'd like to move beyond the Bayesian framework for understanding rational belief. Thus far, we've moved in the direction of greater plausibility by constraining Subjective Bayesianism—making it less permissive. But I actually think that Constrained Subjective Bayesianism is too restrictive to encompass the most rational reactions to certain instances of disagreement. The clearest reason for this is illustrated by the Logic Problem case considered above. Given that the most rational response to that situation is for the agent to remain extremely confident in A, B and C while having only moderate credence in P, the most rational response to this situation will violate probabilistic coherence, and one will not arrive at this response by standard conditionalization. So allowing for the requisite conciliation in this case requires

a move in the direction of relaxing the standard Bayesian assumptions that the most rational response to evidence will always result in a probabilistically coherent set of credences, and be modeled by standard conditionalization on some item of evidence.³¹

We may also make one more move in order to get further away from uniqueness. Constrained Subjective Bayesianism, like pure Subjective Bayesianism, is only permissive interpersonally: in requiring that an agent change beliefs only by conditionalizing, Constrained Subjective Bayesianism imposes a kind of intrapersonal version of uniqueness. So it's worth noticing that even this dimension of uniqueness is not required for conciliationism. In allowing for some non-conditionalization-based belief-changes, one may leave room for a single agent to choose freely among a set of fully rational doxastic alternatives. Clearly, allowing for this is fully compatible with requiring conciliatory responses to the disagreement of others when agents have strong dispute-independent reason to regard them as accuracy-peers.

For ease of reference, let us call the sort of view we've been sketching a Sanely Permissive Account of Rational Credence (SPARC).³² It should be noted that the sort of conciliationism implied by the envisioned forms of SPARC may not quite match the conciliatory views that some would reject on the grounds that they require uniqueness. That will depend on how "peerhood" is cashed out.

If one understood conciliationism as requiring conciliation with apparent *rationality-peers*, then SPARC might well undermine it. Given SPARC's permissiveness, one may well have good reason to believe oneself to be in a relevantly permissive situation: one in which one's own belief and one's friend's belief are both fully rational.³³ Clearly, the *rationality* of one's friend's differing belief does not impugn the rationality of one's own initial belief. Moreover, as noted earlier, most permissive epistemologies will allow for cases in which one sees one's friend's belief as fully rational, yet likely to be inaccurate. So one's friend's disagreement need not impugn the accuracy of one's initial belief either. In such cases, SPARC will allow for disagreement of a rationality-peer that doesn't provide significant reason for conciliation. This is, I think, what people have had in mind when they've thought that reasonable permissive accounts of rationality would undermine conciliatory views of disagreement.

Nevertheless, things look quite different if conciliationism is understood as requiring conciliation with *accuracy-peers*. In that case, conciliationism will be perfectly consistent with, or even entailed by, versions of SPARC. Moreover, many of the sorts of cases people worry about in the disagreement literature—cases in which one disagrees with those who are equally well-informed, intelligent, unbiased, careful and thorough thinkers—are cases that are very plausibly categorized as examples of disagreement by apparent accuracy-peers. After all, well-informedness, intelligence, lack of bias and care in thinking are exactly the sorts of things we believe lead to accurate beliefs.

As noted above, writers have not always been maximally clear about which notion of peerhood they have in mind, and I won't spend time here on textual exegeses. The important point is this: insofar as one is concerned about conciliationism because of its implications for losing rational confidence in controversial beliefs, one's worries

should not be assuaged much at all by arguments against uniqueness. As long as one has good (dispute-independent) reason to think that those who disagree are equally likely to form accurate beliefs on the disputed matter, one may face rational pressure to conciliate, even given a highly permissive account of rationality.

4.4. Conciliation and Doxastic Uniqueness?

The discussion in the last section bears on a recent argument from Stewart Cohen (2013), which connects conciliationism to a weaker sort of uniqueness principle. On Cohen's view, disagreement of rationality-peers generates conciliatory pressure, but it does this *by way of* accuracy considerations. Cohen's idea is that if one has good reason to think that one's equally-informed friend is a rationality-peer, one has no reason to think that one's own belief is more accurate, and that this will motivate conciliation. Cohen argues that his conciliatory view does not require uniqueness. Instead, it requires what Cohen calls "Doxastic Uniqueness":

Doxastic Uniqueness: A subject cannot rationally believe there are two (or more) rational credences for h on e , while rationally holding either (Cohen 2013, p. 101).

According to Doxastic Uniqueness, while there might be two separate agents who rationally have different credences on the same evidence (*contra* uniqueness), an agent cannot rationally hold one of those credences while believing that the other would also be rational.

Cohen does not see commitment to Doxastic Uniqueness as a liability, because he believes that it is true. Nevertheless, it seems to me that if conciliationism in general did require Doxastic Uniqueness, this might offer comfort to critics of conciliationism. One might see Doxastic Uniqueness as intrinsically implausible (the full-belief analogue of it would be violated, for instance, by Fantl and McGrath's account of justified belief). Moreover, it might be argued that Doxastic Uniqueness without full uniqueness was hard to motivate, creating pressure for the conciliationist to embrace uniqueness after all.³⁴

However, as the example of Constrained Subjective Bayesianism shows, strong conciliatory pressure may arise even on views of rationality that violate not only uniqueness, but Doxastic Uniqueness. After all, Constrained Subjective Bayesianism allows an agent to consider other priors—say those priors that weight simplicity more strongly than the agent's own prior does—as perfectly rational, but less likely to be accurate than her own. So there's nothing irrational, on the Constrained Subjective Bayesian account, about an agent having one credence while deeming a different credence perfectly rational. And the same goes, of course, for various versions of SPARC.

So while I do not want to argue against Doxastic Uniqueness, it seems to me that even fairly strong conciliationism does not commit one to it. I should emphasize that the above considerations do not falsify the claim Cohen makes in his article, which concerns a particular version of conciliationism in which epistemic peerhood is understood in terms of rationality, not accuracy. My point here is that one of the insights underlying Cohen's discussion—that the fundamental reason for

conciliating concerns accuracy rather than rationality—actually frees the conciliator from the burden of defending even Doxastic Uniqueness.

5. Accuracy-based Conciliation and Rational Toxicity

I would now like to return to the second interesting feature of conciliationism mentioned at the outset: that it can result in one's being required to adopt beliefs that violate some rational ideal.

It may seem natural that this sort of outcome would result from compromising with the views of others on the basis of evidence that one's own initial take on the evidence is not maximally rational. After all, this sort of evidence may well be misleading, supporting false claims about what reactions to one's initial evidence are rational. If one heeds misleading evidence about what reaction to one's evidence is rational, it may seem unsurprising that one ends up violating some rational ideals.

However, if conciliation is motivated by *accuracy* considerations, and if two equally-informed accuracy-peers may disagree while both being maximally rational, then the tendency of the relevant sort of disagreement to serve as misleading evidence of one's own rational shortcoming is much less obvious. Thus one might wonder whether the sort accuracy-peer-based conciliatory view we are considering will have the same sort of potential for toxicity—whether it will lead to the same sorts of rational pressure to violate rational ideals.

One might, after all, point out that any ordinary evidence one gets that justifies a change in one's belief thereby constitutes evidence against the accuracy of one's initial belief. If I believe we'll eat linguine with clams for dinner, and my partner tells me that we're out of linguine, I should change my belief because it's no longer likely to be true. Now this sort of belief-change certainly does not typically seem to involve violations of rational ideals. And if conciliation with other agents is based on accuracy considerations rather than worries about rationality, one might think that the opinions of others serve just like my partner's testimony about our lack of linguine—ordinary indicators of states of the world, to be accommodated by ordinary rational means. Why should treating other agents as accurate truthometers be any more prone to result in toxicity than treating certain mechanical devices as accurate thermometers in forming beliefs about the temperature?

However, things are not this simple. We can see this right away by considering the Logic Problem case. That case can certainly be taken as one in which conciliation is motivated by accuracy-evaluation of one's friend. And it seems as plain a case of rational toxicity as any.³⁵ So it's clear that toxicity can occur in a case where conciliation is motivated by accuracy-worries about one's initial belief, rather than by rationality-worries. However, the Logic Problem case is also unrepresentative in an important way: It seems to be an assumption of the Logic Problem case that rationality is not permissive with respect to the belief in question: a sort of local uniqueness applies. Thus it's clear from the outset that at least one of the two agents' initial beliefs falls short of maximal rationality. So my friend's disagreement clearly does cast doubt on the rationality of my initial belief. Even if this reason for doubt does not serve as my motivation for conciliation, one might wonder whether its

presence, or the local uniqueness from which it springs, is essential to cases where rational toxicity occurs.

So it's worth asking an additional question: Can rational toxicity occur when disagreement occurs in a permissive case (that is, in a case where more than one opinion can be rational, given one's evidence), and when one does not have reason to think one's original opinion is rationally defective? It seems to me that it can. Let us suppose that a version of SPARC is true. Dinah and Ella are equally-talented meteorologists, who have amassed long and equally impressive track-records of accuracy in weather-prediction. Moreover, let us assume that they've also amassed long and equally impressive track-records of rationality, even when they've disagreed. Perhaps they've been scored by the Epistemology Oracle, who has informed Dinah that she and Ella have almost always reached rational credences, even in cases they've disagreed.

Today, Dinah and Ella are arriving at credences in it raining tomorrow. The chance of rain is increased by the fact that a cold front is approaching. But the front's effect on tomorrow's weather will depend on a number of factors: whether and how its rate of movement changes, whether and how its path alters, and whether it dissolves before it arrives. These factors in turn depend on further factors: the front's speed and direction depend on patterns of wind distribution, and the possibility of the front's dissolving depends on air temperatures high above the surface. The way that these various factors bear on one another, and ultimately on the probability of rain, requires weighting various indicators against one another, assessing the suitability of certain statistical models, and interpreting evidence in accordance with confirmation-theoretic desiderata such as explanatory goodness, simplicity, and fit with the data. In short, this is exactly the sort of case where proponents of permissive epistemologies see room for maximally rational agents to disagree. Suppose that SPARC does indeed allow a certain range of credences to be maximally rational, given the evidence which Dinah and Ella share.³⁶

Suppose that Dinah arrives at .4 credence in rain tomorrow, and then learns that Ella, on the basis of all of the same evidence, has arrived at .6. Dinah, being a committed conciliationist, reduces her credence in rain tomorrow to around .5. However, though the case is a permissive one, suppose that the meteorological evidence they share in fact only permitted credences between .35 and .45. Ella, uncharacteristically, has blundered—perhaps, due to overweighting simplicity, she put insufficient importance on the decreased stability of weather systems in summer, which led to overestimating the directional stability of the cold front and becoming over-confident in rain. In such a case, it may well also happen that Dinah's post-compromise credence violates a rational ideal, by failing to give the meteorological evidence its due.

One might be suspicious here, though, that I'm missing something important about treating people as accuracy-peers: Why isn't the evidence provided by Ella's belief just like the evidence provided by another meteorological instrument? Surely, taking instrument-readings into account doesn't typically involve violating rational ideals! Even if we've stipulated that a credence of .5 would have violated a rational ideal if it were based on Dinah's evidence before learning of Ella's belief, why think

that Dinah's post-compromise belief, which is based on this additional piece of evidence, would also violate a rational ideal?

To answer this question, it is helpful to think (at least schematically) about the sort of ideal that Dinah may be violating. In the Logic Problem case, I was supposing that the rational ideals included probabilistic coherence. But as we saw in thinking about Subjective Bayesianism, it's implausible that this is the only rational constraint. So it's plausible that among the rational ideals are the sort of general confirmation-theoretic constraints on scientific theorizing mentioned above; it is via these principles that the original meteorological data have implications for the rational credence in rain. In this case, it turns out that these implications preclude assigning rain a credence as high as .5.

Now we might ask: why is Ella's opinion not just another data point that changes the probability of rain, according to these very rational principles, when Dinah learns of it? The answer is that the confirmation-theoretic principles, applied to Dinah's total evidence, need not give Ella's opinion much significance. Dinah's background information about Ella includes her history of making accurate predictions on the basis of meteorological data, but also includes other facts about Ella. For one, it includes the fact that her present prediction was made purely on the basis of a particular meteorological data set, and that she has no independent access to rain-relevant facts. (If Dinah thought that Ella might have some independent access to weather-facts—say through her bunions aching—that would be a different matter.) Second, Dinah's information about Ella includes the fact that she is a fallible human who sometimes makes mistakes in interpreting data. Given all this, the confirmation-theoretic constraints applied to Dinah's total evidence may well strongly favor the hypothesis that Ella has made a blunder in coming to such a high credence in rain, so that adding the facts about Ella and her opinion to the meteorological data set has little or no impact on the probability of precipitation.

But if that's right, why can't Dinah dismiss Ella's opinion on exactly these grounds? Why should she conciliate at all? The conciliationist answer to this is that Dinah could dismiss Ella's opinion on this occasion *only via Dinah's own assessment the original data-set*. But Dinah also has strong dispute-independent reason for believing that she and Ella are equally likely to make accurate predictions in cases where their analyses of meteorological data diverge. So although her original opinion about how strongly the data supported rain was one of the rational ones, Dinah cannot rely on her analysis of the data to dismiss Ella's disagreement.³⁷ She must bracket this reasoning in assessing the import of Ella's opinion. And this is what leads her to conciliate with Ella, thereby violating the confirmation-theoretic constraints on her credence in rain.

This point also allows us to see why disagreement evidence is prone to rational toxicity in a way that instrument readings are not. Learning of new instrument readings does not typically cast doubt on one's own assessment of the import of the previous data. Thus in assessing the importance one should place on the new instrument reading, there is no need to bracket one's original reasoning in

accommodating the new bit of evidence, and the mechanism that produces toxicity is not present.³⁸

In the end, the mechanism responsible for the toxicity in the meteorology case is really not different from the mechanism responsible for toxicity in the Logic Problem case. There, the rational ideal of probabilistic coherence mandates that P cannot get less credence than the conjunction of A, B and C—irrespective of my friend’s testimony, and her excellent track record in logic. Probabilistic coherence forbids my giving any credence to the possibility that my friend’s combination of credences—low credence in P and very high very credence in A, B and C—should be expected to be as highly accurate as her credences usually are. Nevertheless, a conciliatory view will say that given our track records in disagreements of this type, I’m rationally precluded from relying on my own judgment about the disputed logical relations to dismiss my friend’s opinion on this occasion. When I bracket that, my accuracy-expectation for my friend’s opinion is high enough to require conciliation. And when I conciliate, I end up violating a rational ideal. In the meteorological case, the range of opinions permitted by the data is of course wider. And it’s less clear what limits the confirmation-theoretic rational constraints impose on credence in rain. But this just makes it even more plausible that Dinah cannot rationally be confident that Ella has blown it this time.

This response suggests one final question one might have about rationality and accuracy-based conciliationism. In the last two cases, it has been highlighted that the main agent’s reason for bracketing some of her own reasoning in assessing the epistemic import of another’s disagreement flows from the irrationality of the agent relying on parts of her own reasoning which the disagreement has called into doubt. Does this not show that despite my setting the examples out in accuracy-peer form, they ultimately rest on disagreement requiring an agent to doubt her own rationality?

One thing that’s clear from our meteorological case is that Dinah’s conciliation need not have anything to do with any doubt on Dinah’s part that *her original belief* was maximally rational. In fact, permissivism precisely makes room for the disagreement of an equally-informed peer to raise doubts about the accuracy of one’s belief without casting doubt on its rationality. (To emphasize this point, we might imagine that the Epistemology Oracle mentioned to Dinah during their chat that Dinah’s original credence is one of the fully rational ones, given the data.) So the toxicity does not stem from the agent’s doubting the rationality of her original belief, which is exactly the sort of doubt that permissive critics of conciliationism have sought to eliminate by criticizing uniqueness.

Is some other sort of doubt on the agent’s part about her own rationality required? To support the requirement to bracket, must the agent worry that her assessment of the friend’s belief might be irrational? I don’t think that even this is the case. Suppose, for example, that Dinah believes in the EHTR, and is quite sure that no assessment she makes of Ella’s belief will engender a harsh vibe. Still, Dinah may bracket her own reasoning about rain in assessing Ella’s reliability, because she worries that her assessment of Ella based on this reasoning could be

inaccurate. And the same point applies even if Dinah believes, more reasonably, that a SPARC-style permissive account of rationality is correct. Dinah need not be fully certain of what precise limits rationality places on rational credence in rain in her situation. But since she believes that rationality is permissive, she may well worry about her assessment of Ella's belief falling short in accuracy, without worrying about it falling short rationally. So it turns out that the rational toxicity produced by disagreement need not derive from an agent's doubts about her own rationality, even in this more indirect way.

6. Conclusion

We've seen that conciliationism can be motivated by two different ways of regarding those with whom one disagrees: as rationality-peers or as accuracy-peers. The former sort of conciliationism sits less easily with denying uniqueness (though quite a bit of conciliation might be called for on rationality grounds, even if uniqueness is false). The latter sort of conciliationism, however, is much less threatened by permissive accounts of rationality.

It turns out that there are two ways in which permissive accounts of rationality can undermine conciliatory principles. The crude way is that a sufficiently permissive epistemology simply puts few constraints on how one reacts to evidence in general, and that can include the evidence provided by the disagreement of others. The less crude way is that permissive epistemologies allow one in certain cases to acknowledge the disagreement of rationality-peers without casting strong doubt on the rationality of one's own initial opinion. But while this can diminish the force of conciliatory pressure based on assessments of rationality-peerhood, it does not similarly undermine conciliatory pressure based on accuracy-peerhood. Given that many of the standard cases of epistemic peerhood are understandable in accuracy-peer terms, a robust conciliationism may find a comfortable home in some highly permissive epistemologies.

Finally, we've seen that while taking conciliationism in this way allows it to flourish in permissive environments, it does not prevent the position from leading to rational toxicity, even in permissive cases when conciliation does not involve an agent's doubts about her own rationality. So the sort of conciliationism that is compatible with strong denials of uniqueness turns out to have both of the features that, for some, make conciliationism alarming: it may preclude having rational confidence about many controversial questions, and it may lead to situations in which the most rational response available to an agent involves violating a rational ideal.

Notes

¹ For comments on earlier drafts and for other helpful discussions of these issues, I'd like to thank Felicia Nimue Ackerman, Nomy Arpaly, Nathan Ballantyne, Stew Cohen, Sophie Horowitz, Matthew Lee, Ben Levinstein, Josh Schechter, Miriam Schoenfield, Jonathan Vogel, and audiences at the Bled Epistemology Conference 2013, the APA Central Division 2013, and the Arché Workshop on Social Epistemology at the University of St. Andrews.

² See, e.g., Feldman (2007), Christensen (2007).

³ See, e.g., Douven (2009), Kelly (2010), Ballantyne and Coffman (2012), Schoenfield (2014).

⁴ See, e.g., Christensen (2007, 2009), Ballantyne and Coffman (2012), and Cohen (2013). White (2013) doesn't take a stand on the entailment issue, but notes that on the face of it, permissivism allows non-conciliatory reactions to disagreement.

⁵ See, e.g., Elga (2007), Kornblith (2010), and Frances (2010).

⁶ Elga (2010) makes a point of noting that his Equal Weight view carries no commitment to uniqueness. Levinstein (ms.) argues that the two issues are orthogonal. See also Lee (2012, 2013), which see the issues as connected, but argue that conciliationism doesn't require uniqueness.

⁷ Why can't I just dismiss my friend's opinion on the basis of my reasoning from A, B and C to P? Conciliatory views require an agent to determine the seriousness with which she takes another's disagreement in a way that's *independent* of the agent's own particular reasoning on the disputed matter. The thought behind this is that insofar as disagreement can raise doubts about the cogency of one's reasoning on a certain matter, it would be question-begging to dismiss those doubts in a way that relied on the very reasoning under dispute.

⁸ The term is from Christensen (2010), which contains an extended development of this idea, arguing that the same phenomenon occurs in the most rational responses to doubts about one's own thinking that are raised by other sorts of evidence—e.g., evidence that one has been administered judgment-distorting drugs. See also Christensen (2007b).

⁹ For arguments at least roughly along these lines, see Titelbaum (forthcoming), and Lasonen-Aarnio (forthcoming).

¹⁰ Or so will I argue. See Schoenfield (2014) for an opposing take on this.

¹¹ Those who support conciliation in cases of peer disagreement naturally say the same for cases of disagreement with epistemic superiors. In fact, the arguments for conciliating with peers seem to me to support at least some conciliation with those one has reason to believe are below the level of one's peers. But I'll suppress these complications for the most part, and just consider apparent peers. I'll also follow the simplifying practice of most of the literature and concentrate on two-person disagreements.

¹² For discussions of peerhood in rationality-based terms, see Feldman (2007), Kelly (2005), Christensen (2007), and Cohen (2013).

¹³ Discussions of peerhood from an accuracy perspective include Elga (2007), White (2009), Enoch (2010), Kelly (2010), Lam (2011) and Levinstein (ms.). When beliefs are thought of in a categorical way, accuracy comes down to truth. Things are not as simple if beliefs are thought of as graded. One can use measures of accuracy for graded beliefs such as the standard Brier score used to assess the accuracy of probabilistic weather forecasts. (In fact, there are different measures of accuracy which have interesting implications for accuracy-based conciliation (see Lam (2011, 2013).) On such a model, a judgment of accuracy-peerhood would come down to assigning the friend's opinion equal expected accuracy (where that assignment is taken to be independent of one's reasoning on the issue under dispute). A related approach to taking the beliefs of others as indicators of what the world is like, without invoking measures of accuracy, is given in Elga (2007). For present purposes, I hope it will suffice if we can think informally about the "accuracy"-based approach as evaluating others' beliefs in terms of their tendency to indicate states of the world, rather than in terms of their conformance to the rules of rationality. And while some measures of accuracy may not fit well with conciliatory views, others (such as Brier score) do, and I'll presuppose some such understanding of accuracy here.

¹⁴ For an early example of this sort of confusion, see Christensen (2007). Christensen's discussions of general motivations for conciliating focus on evaluating whether the other agent forms the (rationally) correct conclusion, given her evidence. But in discussing particular example applications, Christensen often invokes track-records of reaching accurate conclusions. In some cases—as in mathematical calculation—these may coincide. But in others, the notions clearly come apart.

It's also worth noting that some writers use locutions that could be read either way, referring to the other agent's "cognitive capacities" or their being "equally good evaluators of evidence" in relevantly similar situations. For present purposes, it doesn't matter who has had what notion of peerhood in mind; what matters is that we keep the notions separate here. One writer who has clearly distinguished these strains of thought (primarily with reference to the problem of irrelevant influences on belief rather than the problem of disagreement) is Miriam Schoenfield; see her (2014).

Finally, one might worry that on some views of rationality (e.g. some forms of reliabilism), the distinction I'm making here will collapse. (Thanks to Nathan Ballantyne for pointing this out.) However, I think that this will not affect the arguments below, which concern the general relations among conciliationism, uniqueness and rational toxicity.

¹⁵ Roger White (2005) calls this position "Moderate Permissivism."

¹⁶ See Lee (2013) for the point about categorical belief, and Christensen (2007) for the point about credences.

¹⁷ A point along these lines is made in Christensen (2009); see also Ballantyne and Coffman (2011). We might also note that Fantl and McGrath's view is clearly intended to allow cases where agents with the same evidence disagree in the sense that one withholds belief and the other believes, but it does not (if my understanding is correct) permit agents with the same evidence to disagree in the sense of believing *P* and believing $\sim P$. Thus it also serves to illustrate the point of the previous paragraph.

¹⁸ Of course, the accuracy-based approach in no way precludes the possibility that an agent's reason for thinking her friend accurate may include reasons for thinking that her friend responds rationally to evidence. It just makes room for other reasons for expecting a friend's beliefs to be accurate. In what follows, I'll concentrate on agents who have very direct, track-record-style evidence for the accuracy of other agents. But the lessons will apply equally to agents who have other sorts of evidence of another agent's accuracy, be it from information about their rationality, from testimony about the accuracy of the other agent, or from assessments of things like intelligence or carefulness, which may bear on accuracy whether or not it's by way of rationality. Thanks to Ben Levinstein for prompting me to emphasize this.

¹⁹ This point is made clearly in Levinstein (ms.).

²⁰ This problem was pointed out to me by Nathan Ballantyne, in response to a draft of Christensen (2009) which attempted to use the Epistemic Hippie example to show that conciliationism motivated by accuracy considerations did not require uniqueness.

²¹ This is an extremely pure version of the Subjective Bayesian view, which may not have many advocates today. But it's just being used to make a point about certain kinds of accounts of rationality, so this should cause no problem.

²² Note that while Subjective Bayesianism is extremely permissive in one sense, it's also extremely impermissive in another sense. The view precludes what Kelly (2013) calls "intrapersonal slack": it doesn't give *Bruno* any latitude in what to believe, since he's stuck with conditionalizing on his evidence, given his own prior. Douven (2009) and Weintraub (forthcoming) also distinguish between principles that are interpersonally permissive but impose uniqueness intrapersonally, and principles that are permissive even at the intrapersonal level. As these writers point out, it's the more plausible interpersonal sort of permissiveness that seems relevant to the disagreement issue.

²³ Schoenfield (2014) explicitly embraces this consequence of her permissive view, and shows how it enables a permissivist to resist certain attacks. (The central insight is that once a permissivist sees her own epistemic standards as more accuracy-conducive than other perfectly rational standards, she need not see believing in accordance with her standard, rather than other rational standards, as arbitrary.)

²⁴ I suspect that implicitly rejecting this sort of possibility may lie behind some defenses of uniqueness (see Christensen (2007) and White (2005)). White (2013) presses worries about permissiveness that seem to me to be closely related to this point.

²⁵ It is possible for a kind of permissivist to avoid this consequence, by accepting Doxastic Uniqueness (see below).

²⁶ Why "little cause" instead of "no cause at all"? Insofar as Bruno thinks that a disagreeing rationality-peer with his evidence may have a prior close to his own, the disagreement may be some evidence that he's made a performance error in applying his prior to the current situation. I will suppress this complication for now.

²⁷ I should emphasize that I'm simply stipulating that Bruno has a prior that is reasonable *by conciliationist standards*, in order to show how such an agent might be rationally required to conciliate. Other priors are of course possible on the Subjective Bayesian picture—for example, one which would allow Bruno to respond to Callie's disagreement as follows: "Sure, if I bracketed my own reasoning on the present case, I'd have good reason to think Callie equally likely to be right. But as our particular evidence in this case makes clear, Jocko's the culprit, so Callie must be wrong today."

I am also putting aside for now worries over whether a fully conciliatory position can be accommodated within Subjective Bayesianism; I'll examine this issue below.

²⁸ Following the general point made in fn. 18, it may be worth noting that, (given certain background assumptions) Bruno could have reason for thinking that Callie's rationality was an indicator of her accuracy. In such a case, evidence of her rationality could end up motivating conciliation. So the idea here is not that in situations of this type, rationality-information about another agent is always irrelevant. The idea is that its importance in cases like this would come from taking it as indicating the other agent's accuracy, not as indicating that one's original opinion was rationally defective.

²⁹ I am aware that Gandhi's and Bush's lives overlapped for a couple of years. That's how flexible Subjective Bayesianism is.

³⁰ In fact, Elga's Equal Weight View can be seen as a version of this sort of view. And Levinstein follows a somewhat similar pattern in arguing that conciliatory pressure and uniqueness are orthogonal, though he rejects the sort of Independence requirement Elga's account incorporates.

³¹ There are other, more subtle, problems as well with formulating conciliatory views in a strict Bayesian framework. See Lasonen-Aarnio (2013) for extended development of this point.

³² In calling this sort of theory "sanely" permissive, I don't mean to imply that any such theory is correct. Here, I want to prescind from the question of whether such an account is true, to examine the question of what implications the truth of such an account would have for conciliationism.

³³ Kelly (2010) employs this sort of case in arguing that a certain conciliatory view requires uniqueness.

³⁴ For related discussion see White (2005) and Ballantyne and Coffman (2012).

³⁵ This understanding of the Logic Problem case assumes, of course, that our version of SPARC allows in general for the maximally rational response to some situations to involve violating rational ideals. This suggests that it will involve several rational principles or ideals which in some cases compete with one another, having a structure more like Rossian morality than like utilitarianism. One can, of course, reject all such theories, and rule out the phenomenon of rational toxicity from the outset. Our present discussion, however, concerns the question of whether accuracy-based conciliationism can avoid toxicity in a way that rationality-based conciliationism cannot. For that purpose, I will assume that rational toxicity is a possibility. For defenses of this sort of view, see the references in note 8.

³⁶ Thanks to Felicia Nimue Ackerman for help in devising this meteorological example.

³⁷ Of course, if Ella's interpretation of the data struck Dinah as completely absurd, she might have reason to think that Ella was subject to some unusual malfunction. (Such a case might arise if Ella reported having credence 1 in rain.) In such cases, Dinah might not be required to conciliate much, or at all, with Ella's expressed opinion. (See Elga (2007) or Christensen (2007, 2011) for explanations of how conciliatory theories allow steadfast responses in such cases; the rough idea is that those cases allow one to downgrade the expected accuracy of the other person in a way that does not rely on one's own reasoning about the matter under dispute.) But in our present case, Ella's opinion does not seem crazy at all, so Dinah lacks significant dispute-independent reason for thinking that Ella has blown it.

³⁸ It should be pointed out that this difference between other people's disagreement and instrument-readings is not absolute. There are certainly conceivable cases where instrument-readings can support the sort of self-doubt that motivates the bracketing involved in Dinah's case. But typical instances of learning instrument-readings are not like this, and so don't lead to significant rational toxicity.

References

- Ballantyne, N. and E. J. Coffman (2011), "Uniqueness, Evidence and Rationality," *Philosophers' Imprint* 11: 1–13.
- . (2012), "Conciliationism and Uniqueness," *Australasian Journal of Philosophy* 90: 657–670.
- Christensen, D. (2007), "Epistemology of Disagreement: the Good News," *Philosophical Review* 116: 187–217.
- . (2007b), "Does Murphy's Law Apply in Epistemology? Self-Doubt and Rational Ideals," *Oxford Studies in Epistemology* 2: 3–31.

- . (2009), “Disagreement as Evidence: The Epistemology of Controversy,” *Philosophy Compass* 4: pp. 756–67.
- . (2010), “Higher-Order Evidence,” *Philosophy and Phenomenological Research* 81: 185–215.
- . (2011), “Disagreement, Question-Begging and Epistemic Self-Criticism”, *Philosophers’ Imprint* 11, no. 6.
- Cohen, S. (2013), “A Defense of the (Almost) Equal Weight View,” in D. Christensen and J. Lackey, eds., *The Epistemology of Disagreement: New Essays* (Oxford: Oxford University Press).
- Douven, I. (2009), “Uniqueness Revisited,” *American Philosophical Quarterly* 46, 347–361.
- Elga, A. (2007), “Reflection and Disagreement,” *Noûs* 41: 478–502.
- Enoch, D. (2010), “Not Just a Truthometer: Taking Oneself Seriously (but not Too Seriously) in Cases of Peer Disagreement,” *Mind* 119: 953–997.
- Fantl, J. and M. McGrath (2002), “Evidence, Pragmatics and Justification,” *Philosophical Review* 111: 67–94.
- Feldman, R. (2007): “Reasonable Religious Disagreement,” in L.M. Antony, ed., *Philosophers Without Gods: Meditations on Atheism and the Secular Life*, (Oxford: Oxford University Press) 194–214.
- Feldman, R. and T. Warfield, eds. (2010), *Disagreement* (New York: Oxford University Press).
- Frances, B. (2010), “The Reflective Epistemic Renegade,” *Philosophy and Phenomenological Research* 81, 2: 419–463.
- Kelly, T. (2005) “The Epistemic Significance of Disagreement,” *Oxford Studies in Epistemology* 1: 167–96.
- . (2010), “Peer Disagreement and Higher-Order Evidence,” in Feldman and Warfield eds., *Disagreement* (New York: Oxford University Press).
- . (2013), “How to be an Epistemic Permissivist,” in M. Steup, J. Turri and E. Sosa, eds., *Contemporary Debates in Epistemology* (Malden, MA: Wiley-Blackwell).
- Lam, B. (2011), “On the Rationality of Belief-invariance in Light of Peer Disagreement,” *Philosophical Review* 120: 207–245.
- . (2013), “Calibrated Probabilities and the Epistemology of Disagreement,” *Synthese* 190: 1079–1098.
- Lasonen-Aarnio, M. (forthcoming), “Higher-Order Evidence and the Limits of Defeat,” *Philosophy and Phenomenological Research*.
- . (2013), “Disagreement and Evidential Attenuation,” *Noûs* 47: 767–794.
- Lee, M. (2012), “The Role of the Uniqueness Thesis in the Epistemology of Disagreement,” in C. Jäger and W. Löffler, eds., *Proceedings of the 34th International Wittgenstein Symposium: Epistemology: Contexts, Values, Disagreement* (Ontos Verlag).
- . (2013), “Conciliationism without Uniqueness,” *Grazer Philosophische Studien* 88: 161–188.
- Levinstein, B. (ms.) “What your Credence Tells me about P,” chapter 4 of *Accuracy as Epistemic Utility*, Ph.D. dissertation at Rutgers University.
- Schoenfeld, M. (2014), “Permission to Believe: Why Permissivism is True and What it Tells us about Irrelevant Influences on Belief,” *Noûs* 48: 193–218.
- Titelbaum, M. (forthcoming), “Rationality’s Fixed Point (Or: In Defense of Right Reason),” *Oxford Studies in Epistemology*.
- Weintraub, R. (forthcoming), “Can Steadfast Peer Disagreement be Rational?” *Philosophical Quarterly*.
- White, R., (2005), “Epistemic Permissiveness,” *Philosophical Perspectives* 19: 445–459.
- . (2009), “On Treating Oneself and Others as Thermometers,” *Episteme* 6: 233–250.
- . (2013), “Evidence Cannot be Permissive,” in M. Steup, J. Turri and E. Sosa, eds., *Contemporary Debates in Epistemology* (Malden, MA: Wiley-Blackwell).