

DISAGREEMENT AND EASY BOOTSTRAPPING

EYAL TAL

eyaltal@email.arizona.edu

ABSTRACT

Should conciliating with disagreeing peers be considered sufficient for reaching rational beliefs? Thomas Kelly argues that when taken this way, Conciliationism lets those who enter into a disagreement with an irrational belief reach a rational belief all too easily. Three kinds of responses defending Conciliationism are found in the literature. One response has it that conciliation is required only of agents who have a rational belief as they enter into a disagreement. This response yields a requirement that no one should follow. If the need to conciliate applies only to already rational agents, then an agent must conciliate only when her peer is the one irrational. A second response views conciliation as merely necessary for having a rational belief. This alone does little to address the central question of what is rational to believe when facing a disagreeing peer. Attempts to develop the response either reduce to the first response, or deem necessary an unnecessary doxastic revision, or imply that rational dilemmas obtain in cases where intuitively there are none. A third response tells us to weigh what our pre-disagreement evidence supports against the evidence from the disagreement itself. This invites epistemic akrasia.

I. INTRODUCTION

According to Conciliationism, one must revise one's view upon meeting an epistemic peer who has a different view.¹ The further a peer's view is from one's own, the greater the revision required. On a graded doxastic framework, many understand Conciliationism as requiring both parties to revise their respective credences in the direction of their peer's credence.

Thomas Kelly (2010, 2014) worries that Conciliationism makes it too easy to have fully rational doxastic attitudes. The argument follows a simple recipe. First, take an agent who has an irrational credence in a proposition p . Then, introduce that agent to a peer who disagrees. Finally, per Conciliationism, have the agent conciliate to some extent. And there you have it: the agent is now rational in having the credence that she does in p . The problem is most vivid when both agents start out with irrational credences, and end up with rational ones by conciliating with each other. Let us follow Kelly in referring to this as the Easy Bootstrapping argument (EB).² EB shows that an intuitive

1 Many authors, at some point or another, have defended versions of this view. A representative sample includes Christensen (2007), Elga (2007), Feldman (2007), and Cohen (2013).

2 Weatherson (Ms) names this phenomenon “epistemic laundering.”

understanding of Conciliationism, on which the rationally required credence is always the one we reach by conciliating from our initial credence, appears to allow for unacceptably easy formation of rational attitudes.³ Call this intuitive understanding Naïve Conciliationism:

Naïve Conciliationism: If one has credence c_1 in p , and one encounters a peer who has a different credence c_2 in p , then rationality requires that one conciliate to an intermediate credence c_3 .

Attempts to defend Conciliationism from a bootstrapping charge fit into a few categories.⁴ These attempts either (1) restrict the need to conciliate only to agents who have a rational credence as they enter into a disagreement, or (2) take Conciliationism to be a necessary constraint on rational credences rather than a sufficient one, or (3) require that we weigh what our pre-disagreement evidence supports against the evidence that the disagreement brings. I argue that the resulting conciliatory views are all unpromising. Restricting the need to conciliate only to agents who have rational credences ensures that Conciliationism is entirely concerned with misleading those agents whose views are rational. This gives the exact opposite outcome to the one we were hoping to get out of a conciliatory strategy, namely, to help mitigate the mistakes that we make. As such, there are reasons to think that if this restricted version is what Conciliationism amounts to, no one should follow its directives. Alternatively, understanding Conciliationism as a mere necessary condition for having a rational credence is not enough, as it fails to tell us what rationality requires in cases of disagreement. Importantly, natural ways of fleshing out the corresponding requirement have unwanted implications: either they deem necessary a doxastic move that is sometimes unnecessary, or they effectively reduce to the restricting response, or they give rise to rational dilemmas in cases where there do not seem to be any. Lastly, requiring that agents weigh what their pre-disagreement evidence supports against the evidence introduced by the disagreement makes for a rule that we may not rationally believe we followed. Consequently, the rule cements the very kind of doxastic tension that Conciliationism was called to resolve – where we hold a belief despite good reason for thinking that the belief is irrational. An attempt to bite this bullet leads to a contradiction.

2. CONCILIATIONISM AS RESTRICTED TO THE INITIALLY RATIONAL

While discussing a paradigmatic version of Conciliationism – the Equal Weight View – Brian Weatherson (Ms) proposes we restrict the view in order to sidestep EB:⁵

-
- 3 There is a background assumption here that if a credence is rationally required, then having it (in the proper way) is both necessary and sufficient in order to have a rational doxastic attitude. This assumption will be called into question later on (in section 3), but note that it is not unreasonable. A credence that is sufficient but not necessary for believing rationally would not be genuinely required. A credence that is necessary but not sufficient for believing rationally would be irrational to have, which would make rationality require irrationality.
 - 4 EB is an application of the infamous *detaching problem* to Naïve Conciliationism. As a result, some of the possible solutions resemble familiar ones in the normative detaching debate. See Broome (1999), Finlay (2010).
 - 5 Following Elga (2007), Weatherson takes the Equal Weight View to require that we give equal weight to each participant's stance in a peer disagreement.

Consider an agent who makes an irrational judgment. And assume her friend, who she knows to be a peer, makes the same irrational judgment. What does the Equal Weight View say she should do? It should be bad for it to say that she should regard her and her friend as equally likely to be right, so she should keep this judgment. After all, it was irrational! There are a couple of moves the friend of the Equal Weight View can make at this point. But I think the simplest one will be to put some kind of restriction on Equal Weight ... to agents who have initially made rational judgments ... (Weatherson *Ms*: 5)

Weatherson suggests that we understand Conciliationism as requiring conciliation with a peer only when we have a rational credence. With this restriction, when we start out irrational, we are neither required to conciliate nor guaranteed to become rational by conciliating.⁶ So the restriction successfully blocks EB.

A conciliatory rule based on Weatherson's suggestion would look something like the following:

Restricted Conciliationism: If one rationally has credence c_1 in p , and one encounters a peer who has a different credence c_2 in p , then rationality requires one to conciliate to an intermediate credence c_3 .

Restricted Conciliationism allows agents to reach a rational credence by conciliating, but only when those agents have a rational credence in the first place. Since the rule's requirement does not bind agents who enter into a disagreement with an irrational credence, such agents would not necessarily reach a rational credence by conciliating.⁷

Soon after raising EB, Kelly (2010) argues against restricting Conciliationism as a solution:

If the phenomenon of peer disagreement requires you to split the difference with my unreasonable opinion, why should I be spared having to split the difference with your reasonable opinion simply in virtue of having botched the evidence in the first place? Whatever normative pressure is created by the phenomenon of peer disagreement, surely one does not immunize oneself against that

6 Here and throughout, let "agents who have a rational credence/belief" mean the same as "rational agents." By these I mean to refer to agents whose relevant doxastic attitudes are rationally had, i.e., agents who are *doxastically* justified. When I mean to discuss *propositional* justification, I talk of the credences/beliefs that the evidence supports or justifies.

7 Cohen (2013) is sympathetic to such a restriction in response to bootstrapping worries, but also entertains a way for the initially irrational to conciliate to a fully rational credence. Cohen suggests that if the initially irrational agent reflects on her general rationality, her irrational credence could gain enough merit so as to justify conciliation from that credence. By stipulation, the agent should give her peer's view significant evidential weight, and both agents enjoy equal epistemic status. So, the thought goes, the agent should give her own view significant weight as well.

There are places to resist this suggestion. For example, Kelly (2014) argues that it would allow for "even easier, more implausible bootstrapping." The idea is that if rationality requires us to spread the evidential weight across the views of all peers involved, then when only a single peer is involved (i.e., pre-disagreement), her view would get all the evidential weight. In other words, whatever view the agent happens to form before encountering disagreement would be (at least propositionally) justified. Whether or not this is enough to reject Cohen's suggestion, it seems to show that the conciliatory move itself would be required only when the agent is at a rational credence. Thus, Cohen's approach looks like a version of Restricted Conciliationism rather than an alternative to it.

pressure simply in virtue of having beliefs that are not adequately supported by one's evidence. (Kelly 2010: 128)

But the restricting response to EB is not so easily dismissed. The response is not committed to disagreement putting no pressure on irrational agents. Rather, it is silent about what disagreement means for irrational agents, and speaks only to what rational agents should do to maintain their rationality. One could accept the restricted version of Conciliationism and also maintain that disagreement puts additional pressure on irrational agents to revise, beyond the pressure that such agents are already under from their pre-disagreement evidence. Nevertheless, there is something in the neighborhood of Kelly's argument against restricting Conciliationism that proponents of the restriction should worry about.

To get a feel for the problem with Restricted Conciliationism, consider a safety requirement to wear a helmet when riding a bike. Presumably, the only ones who benefit from wearing a helmet are those who will suffer an accident. Since we do not know in advance who will suffer an accident, we think that everyone should wear one. Suppose that someone suggested that the requirement to wear a helmet applied only to those who were *not* going to suffer an accident. Such a suggestion would be absurd. It would get the motivation for wearing a helmet precisely backwards. If we think that those who will suffer an accident need not wear a helmet, it would be hard to explain why anyone else should. So if we want to avoid the result that no one should wear a helmet, we should reject the suggestion. Analogously, if there is any motivation for us to conciliate with our disagreeing peers, it comes from those cases where we are the ones irrational. But Restricted Conciliationism, oddly, requires conciliation only of rational agents.

Conciliationism appears to go awry when it excludes the initially irrational from having to conciliate. The thought should come naturally to those who share David Christensen's (2007) view regarding the good news about the conciliatory strategy:

The fact of disagreement is old, but bad, news; it is bad because it indicates the relatively benighted conditions under which we work. But adjusting our beliefs in the direction of those peers with whom we disagree should be welcomed as a valuable strategy for coping with our known infirmities. (Christensen 2007: 216)

If only the initially rational must conciliate, the news that Conciliationism brings is not good at all. Rather than help us cope with our known infirmities, Restricted Conciliationism stays silent about what we should do when those infirmities lead us astray, while ensuring that our peers' infirmities lead us astray too. If Conciliationism is to be any kind of good news, it cannot reserve the need to conciliate exclusively for the initially rational. The motivation for such a view is hopelessly missing. Call this accusation No Motivation.

Like Kelly's argument against Restricted Conciliationism, No Motivation also focuses on the fact that the view leaves out irrational agents. However, unlike the former, No Motivation is not concerned with whether Restricted Conciliationism lets irrational agents survive disagreement unaffected. Rather, the concern is the irrationality of our disagreeing interlocutors whenever the rule applies to us. Thus, No Motivation seems to rest on the claim that in a peer disagreement, if one peer believes rationally, the other does not:

Someone's Irrational: When two peers disagree, at least one of them believes irrationally.

I discuss Someone's Irrational soon. But first, note an unusual implication of Restricted Conciliationism. Suppose that we come to know that Restricted Conciliationism is a genuine rule of rationality, and we wonder whether we followed it correctly in a given disagreement. If we may rationally believe that we followed it correctly, we may also rationally believe that our peer's credence was the irrational one – for the former entails the latter on the assumption of Someone's Irrational. At that point we would find ourselves in an unstable doxastic position; believing that we changed our view due to a peer's irrational view. We would then have to roll back whatever conciliatory revision we made due to Restricted Conciliationism. So, we get the bizarre result that we cannot rationally believe that we responded to disagreement correctly while maintaining the conciliated view. The rule is inconsistent with rationally believing that we followed it.⁸

The result is more than a quirk of Restricted Conciliationism. It ought to make us wonder who should maintain the rule's recommended credence, if not agents who rationally believe that they followed the rule correctly. Agents who disbelieve or suspend judgment about whether they followed the rule correctly would make for poor candidates.⁹ A major selling point for conciliatory views is that they purport to offer a way to resolve the seemingly irrational tension in having a belief and at the same time having serious justified doubts that the belief is rational. Indeed, these are the doubts that disagreeing peers cause. But agents who follow Restricted Conciliationism and suspend about whether they did so correctly would manifest this very tension, as they would be quite uncertain that their own belief is rational. Agents who follow the rule and disbelieve that they did would fare even worse in this regard. So, since Restricted Conciliationism is inconsistent with a rational belief that we followed it correctly, the rule would seem to secure the tension that it purports to resolve.

Another option is that agents who never consider whether they followed the rule correctly are the only ones who should maintain the rule's recommended credence. This option runs into trouble as well. We could encourage these agents to be more reflective, and thereby undermine their following the rule. It is implausible that the authority of a genuine rule of rationality is contingent on how frequently people consider whether their beliefs are rational.¹⁰

Now, why think that Someone's Irrational is true? The thought that disagreement puts pressure on agents to move away from their initial credences is more compelling when

8 This kind of inconsistency does not arise with other, similarly restricted rules. Suppose, for instance, that we wonder whether something like Modus Ponens is a rule of rationality. No one should think that if we irrationally believe p and also irrationally believe that if p then q , it would be rational for us to believe q (see Harman 1986). But a restricted version of Modus Ponens to those who rationally believe p and if p then q gets the right result, and with no worry of inconsistency. We could come to rationally believe that we followed the rule correctly while maintaining its recommended belief q .

9 Others have argued that that both disbelief and suspension about whether we rationally believe p are inconsistent with believing p . Horowitz (2014) shows that akratic belief combinations of the sort p , but believing p is irrational “license patently bad reasoning and irrational action.” Huemer (2011) and Smithies (2012) take suspension about p to license judgments like *it may or may not be the case that p* , and *it is an open question whether p* . It seems arbitrary to believe p while also believing that *it is an open question whether we are justified in believing p* . I return to this matter in section 4 where a similar issue comes up.

10 A similar point applies with respect to agents who do not know that Restricted Conciliationism is correct. We could inform those agents that it is, and thereby undermine their following it.

disagreement entails that one of the parties involved believes irrationally.¹¹ Paradigm examples used to promote conciliatory views – like Christensen’s (2007) restaurant case – are ones where we know that someone made an error in reasoning. So it would be an interesting upshot of No Motivation if it forced conciliationists to reject Someone’s Irrational, because doing so would eliminate much of Conciliationism’s intuitive appeal.

Someone’s Irrational also follows if we accept a conception of ‘peer’ according to which peers share their evidence, and also accept the Uniqueness Thesis:

Uniqueness: Given one’s total evidence, there is a unique rational doxastic attitude that one can take to any proposition.¹²

Kelly (2010) has argued that Conciliationism is committed to Uniqueness.¹³ But while rejecting Uniqueness would be enough to resist Someone’s Irrational, it would not be enough to resist No Motivation. If Restricted Conciliationism aims to capture what we should do in all possible disagreements, it must be that No Motivation never undermines its recommendation. Yet, whenever we would justifiably believe that one of the disagreeing peers is irrational, we could use No Motivation as grounds for dismissing the recommendation that Restricted Conciliationism makes. So the kind of rejection of Uniqueness required to fend off No Motivation would have to deny that we are ever justified in believing that one of the disagreeing parties is irrational. Of course, disagreements can arise over very many bodies of evidence, and the opinions involved can vary to a great extent. If we are never justified in believing that one of the disagreeing parties is irrational, it must be that permissive situations are both widespread and highly slack.¹⁴ So the brand of Permissivism required for blocking No Motivation, if we wish Restricted Conciliationism to apply to all disagreements, is quite a radical one. It must maintain that very many bodies of evidence are permissive, and that they allow for widely divergent doxastic attitudes.¹⁵

Another way to reject Someone’s Irrational is to deny that peers share all of their relevant evidence. This would allow both disagreeing peers to be rational in cases where their respective evidence differs. Even so, such a rejection of Someone’s Irrational would not undermine No Motivation in cases where the disagreement persists despite explicit sharing of the evidence. So No Motivation is immune to the objection in cases of full disclosure.¹⁶ This defense of Someone’s Irrational works best on the assumption that an agent’s evidence

11 It would be enough for present purposes to take disagreement to suggest that at least one of the agents involved made an error in reasoning (rather than believes irrationally).

12 I borrow this formulation from White (2005).

13 See Kelly (2010: 120–1). Ballantyne and Coffman (2012) discuss the matter in depth.

14 By ‘permissive situations’ I mean a situation in which the evidence permits more than one maximally rational belief. By ‘Permissivism’ I mean the denial of Uniqueness.

15 Cohen’s (2013) suggestion (see footnote 7) appears to afford a kind of widespread Permissivism that undermines Someone’s Irrational. On that suggestion, the track record of generally rational agents seems to (at least propositionally) justify the beliefs that they form. This would make it hard for agents to have irrational beliefs when they are generally rational, and similarly hard to rationally believe that either peer has an irrational belief. Nonetheless, there are cases in which the disagreeing peers should be confident that one of them indeed has an irrational belief. For example, we could inform the agents that one of them has ingested a reason-distorting drug. That would ensure that the agents could not appeal to their track record success on the particular occasion.

16 Feldman (2006) introduces this terminology.

is entirely sharable. But even if our evidence is not always entirely sharable, on many occasions it is. If Restricted Conciliationism is to apply as a rule for belief revision in all disagreement, it must apply when we know that we share all of our evidence with our peers.

It is possible to resist No Motivation by further restricting Conciliationism, so as to guarantee absence of justification for thinking that some party in the disagreement is irrational. Conciliating may thus be taken as required only of initially rational agents who lack justification for believing that one of the disagreeing peers is irrational. Such a rule would be largely incomplete and suspiciously ad hoc. It is hard enough to grant that a rule regarding the correct response to disagreement can leave out all initially irrational agents, but that restriction at least has a principled reason behind it. As Christensen (2011) notes, there are many ways to be irrational, so expecting one rule to account for all possible mistakes may be expecting too much.¹⁷ However, it would be much harder to justify a restriction that leaves out cases of agents who know they are in a disagreement where someone is irrational. Doing so would seem to limit the application of Restricted Conciliationism for no other reason than to strategically avoid No Motivation, while leaving unanswered the interesting question of how we should respond to disagreements when we know that someone is irrational.

The attempt to avoid EB by restricting Conciliationism only to the initially rational does not look promising. Some of the other attempts to avoid EB end up sharing a lot with Restricted Conciliationism, including its problems. Others still face problems of their own. I now turn to these attempts.

3. CONCILIATION AS MERELY NECESSARY

A simple way around EB involves construing Conciliationism as a view about what is necessary for having a rational credence. EB does not gain traction unless we treat conciliation as sufficient for having a rational credence. Kelly (2014) names the construal of Conciliationism that exposes it to EB “Strong Conciliationism”:

Strong Conciliationism: An agent’s holding the conciliated opinion is both sufficient as well as necessary for her opinion to be fully reasonable or justified.¹⁸

On strong conciliatory views, conciliating with a disagreeing peer is required of all agents, and doing so suffices for reaching rational credences. If instead we interpret Conciliationism as imposing just a necessary constraint, conciliating would not guarantee a rational credence. Kelly names such an interpretation “Weak Conciliationism”:

Weak Conciliationism: A perfectly rational agent would not hold some opinion that is out of step with the conciliated opinion.

Weak Conciliationism implies that conciliating with a disagreeing peer is necessary for holding a perfectly rational credence. Since it is not committed to conciliation being sufficient for holding a rational credence, it does not enable easy bootstrapping.

¹⁷ Christensen (2011: 4).

¹⁸ Notice that dropping the necessity constraint would still leave Strong Conciliationism exposed to EB.

On its own, Weak Conciliationism provides little insight into the epistemic significance of peer disagreement. Merely saying that we cannot be perfectly rational as long as we are out of step with the conciliated opinion does not tell us what rationality requires in peer disagreement situations. But just what we want to know is what rationality requires in peer disagreement situations.¹⁹ It is unclear what particular conciliatory requirement follows from Weak Conciliationism, and what other rational requirements are relevant in this context. In fact, Weak Conciliationism does not even tell us whether the perfectly rational agent ever conciliates with her opposition, or whether she ever faces peer disagreement.²⁰ So in the absence of a more detailed story, the epistemic significance of disagreement remains largely mysterious on this view.

Indeed, if Weak Conciliationism is the most we can say about the epistemic significance of disagreement then that would be an interesting result (albeit a disappointing one). That said, some conciliatory accounts in line with the thought that conciliation is only necessary do say more. So, to give the thought due consideration, we should look into possible conciliatory accounts in its spirit.

One tempting way of developing the thought tells us how disagreement affects what is necessary for us to believe rationally:

Necessary Conciliationism: If one has credence c_1 in p (given evidence E), and one encounters a peer who has a different credence c_2 in p (given E), then it is necessary for believing rationally that one conciliate to an intermediate credence c_3 .

According to Necessary Conciliationism, it is necessary that all agents conciliate upon meeting a disagreeing peer, if they are to have a rational credence. Since it is only a necessary condition for believing rationally, meeting it does not guarantee that all who comply end up with rational credences. For the view to avoid EB it must be that the initially irrational would remain irrational even if they satisfy the condition.

However, the conciliatory move that Necessary Conciliationism claims to be necessary sometimes looks entirely unnecessary. Consider the following case:

Fixed Mistakes

Fixit incorrectly evaluates her evidence E , and irrationally forms credence c_1 in p . Fixit's peer correctly evaluates E , and rationally forms the justified credence c_2 in p . Once the disagreement is revealed, Fixit considers her total evidence, correctly judges that it justifies c_2 , and forms c_2 in p based on her total evidence.

In Fixed Mistakes, Fixit finds her way to a perfectly rational credence. She arrives at the credence justified by her total evidence – the pre-disagreement evidence and the new evidence of her peer's credence – and moreover, she does it by reasoning correctly from that

19 For example, Adam Elga (2007) asks: "How should you take into account the opinions of an advisor?" David Enoch (2010) asks: "How should you update your (degrees of) belief about a proposition when you find out that someone else – as reliable as you are in these matters – disagrees with you about its truth value?" Other authors, including Christensen (2009), present the question similarly.

20 One issue here is that it is unclear what the idealization of the agent amounts to exactly. Perhaps the best way to conceive of the agent is as one who always believes rationally, but also lacks sufficient evidence to believe that she does so. Whether this stipulation is coherent is debatable. See Titelbaum (2015) for reasons to doubt it.

total evidence. At no point is it necessary for Fixit to conciliate in the way that Necessary Conciliationism demands. So, contra Necessary Conciliationism, conciliating from our initial credence is not always necessary for arriving at a rational credence.²¹

Two responses to the case suggest themselves, and both rely on the thought that on her way to c_2 Fixit must be changing her credence in p given the pre-disagreement evidence E alone. If this is right, then one way to defend Necessary Conciliationism from the counterexample would be to have the rule apply only to agents who have reached their final verdicts about p given E . Alternatively, we may say that the counterexample misses its mark because as soon as Fixit fixes her irrational credence in p given E , the credence that Necessary Conciliationism recommends changes as well. If upon considering the total evidence Fixit comes to have c_2 in p given E , the rule would recommend that she stay at c_2 . Thus, the case is not a genuine counterexample to the rule.²²

The details of Fixed Mistakes leave it open whether, on her way to c_2 , Fixit changes her credence in p based on the pre-disagreement evidence E alone. If it is possible for her to rationally have c_2 in p based on the total evidence and without forming a new credence in p given E alone, then the counterexample is safe from the two responses. But whether or not that is possible, it would be wrong to restrict Necessary Conciliationism to agents who have reached their final verdicts about p given E alone. Not only is it unclear what reaching such a final verdict amounts to, a conciliatory account should have something to tell agents who may one day change their credence in p given E alone.

The alternative response says that Necessary Conciliationism changes its recommendation as soon as Fixit changes her credence in p given E alone. On this reading of Necessary Conciliationism, it is necessary for believing rationally that we conciliate from whatever credence we have in p given E alone, rather than from the credence in p that we have when we enter into a disagreement. This reading does not escape a more detailed version of Fixed Mistakes. Fixit need not form a new credence in p given E in order to end up fully rational. Fixit's reasoning could proceed via a higher-order judgment about what E justifies. She could rationally judge that E justifies her peer's credence c_2 , and based on that plus the fact that the peer is at c_2 move to c_2 directly.²³

Another way to understand the idea that it is necessary to conciliate from whatever credence we have in p given E is to take the relevant requirement to be wide-scope. I consider such a view next.

3.1 *The wide-scope response*

Christensen (2011) offers a possible response to EB that seems to fit well with the thought that conciliating is only necessary for believing rationally:

-
- 21 This objection may join Kelly's EB in casting doubt on the naive interpretation of Conciliationism. Both the counterexample here and Kelly's bootstrapping accusation trade on the intuition that we should have the right idea of what credence the pre-disagreement evidence justifies, before it is the case that rationality requires us to do anything with that credence (be that requirement necessary or sufficient for a rational belief). Kelly's point targets the sufficiency claim, whereas Fixed Mistakes could be used against the necessity claim.
- 22 I thank an anonymous *Episteme* reviewer for this point.
- 23 It is possible to defend Necessary Conciliationism from Fixed Mistakes by restricting the rule to agents who do not doubt the rationality of their own credence. That would not only be ad hoc, but would also leave Conciliationism silent about what agents who have such higher-order doubts should believe. This would be an odd omission, since disagreeing peers make us have such doubts.

Conciliationism tells us what the proper response is to one particular kind of evidence. . . . But having taken correct account of one bit of evidence cannot be equivalent to having beliefs that are (even propositionally) rational, all things considered. If one starts out by botching things epistemically, and then takes correct account of one bit of evidence, it's unlikely that one will end up with fully rational beliefs. And it would surely be asking too much of a principle describing the correct response to peer disagreement to demand that it include a complete recipe for undoing every epistemic mistake one might be making . . . (Christensen 2011: 4)

If Conciliationism is silent about what credences agents should have, then EB appears to be misguided. On this thought, Conciliationism is merely a view about the proper response to one piece of evidence, or a view about the amount of revision called for by the disagreement on its own, and not a view about how to find an all-things-considered rational belief in any given disagreement.

As with Weak Conciliationism, it is not clear what conciliatory requirement follows when we take conciliating to merely be the correct response to one bit of evidence. Cohen (2013) suggests that Christensen may have a wide-scope requirement in mind:

Wide-Scope Conciliationism: One ought to see to it that if one has credence c_1 in p (given evidence E), and c_1 differs from a peer's credence c_2 in p (given E), then one conciliate to an intermediate credence c_3 .

Having the entire conditional in the scope of the normative operator allows for two ways of complying with the requirement. We could comply with Wide-Scope Conciliationism either by falsifying its antecedent or by satisfying its consequent. Conciliating to c_3 would make the consequent true, while any change to our credence in p (given E) would make the antecedent false.²⁴ Note that any such change that is not a change all the way to the peer's credence c_2 would invite a new instance of the rule, only with different values for c_1 and c_3 . So to comply with Wide-Scope Conciliationism once and for all, at some point we would have to either conciliate or conform to our peer's view. This makes three routes of belief-revision available to us if we are to comply with Wide-Scope Conciliationism in a way that prevents its recurrence. We could conciliate from our credence, or revise our credence in p (given E) so as to conform to our peer's credence, or revise our credence in p (given E) and then conciliate from that new credence.

The wide-scope reading of Conciliationism appears to fit Christensen's description of agents who conciliate as taking correct account of one bit of evidence. If an initially irrational agent conciliates in the face of disagreement, then by doing so she successfully avoids violating a wide-scope requirement. The view also doubly avoids a bootstrapping charge. Complying with the requirement does not guarantee that those who conciliate reach a rational credence, nor does it tell agents who have irrational credences to conciliate. In stating a mere necessary condition for believing rationally, and in staying silent about which way of complying with it is rationally required, Wide-Scope Conciliationism leaves it for other rational requirements to determine what the all-things-considered rational response is to each disagreement.

²⁴ Since the discussion here concerns the rational response to disagreement, we should leave out more creative ways of falsifying the antecedent, like killing our peers.

Yet in its silence about which way of complying with the rule is rationally required in which situations, Wide-Scope Conciliationism leaves obscure how we should respond to a disagreeing peer. It does not tell us who, if anyone, should satisfy the consequent and conciliate, and who should do something else.²⁵ Plausibly, when an agent is the irrational party, the only rational way for her to satisfy Wide-Scope Conciliationism would be to do something other than conciliate from her irrational credence. Any account that includes Wide-Scope Conciliationism must grant this, if the account is to avoid EB. But this also leaves agents who have a rational credence as the only candidates for those who are rationally required to comply with rule by conciliating. So it seems that on a reasonable understanding of how Wide-Scope Conciliationism fits with other rational requirements, it is still only those with rational credences who must conciliate when facing a disagreeing peer. For this reason, here too we could resist the rule with No Motivation and the related inconsistency in rationally believing that we complied with the rule as we should have.²⁶

Wide-Scope Conciliationism is not concerned with how other rational requirements constrain the way that we should comply with it. But when we consider those other requirements, and remember that EB is a looming threat, we see that Wide-Scope Conciliationism could only fit into a picture that reserves the need to conciliate exclusively to those who have a rational credence. The wide-scope interpretation of Conciliationism as merely necessary effectively reduces to the restricted version of Conciliationism.

The fact that we may come to realize past mistakes seems to cause trouble for a narrow-scope interpretation of Conciliationism as necessary. Going wide-scope appears to lead to a previously discussed dead end. Another possibility is to stick to a narrow-scope conciliatory view, and try to make it fit with the rationality of fixing our irrational credences. The next two responses to EB are such attempts.

3.2 *The dilemma response*

When discussing disagreement about disagreement, Christensen (2013) concludes that in certain situations epistemic requirements conflict:

the fact that there is this tension among our epistemic ideals need not mean that any of them is incorrect. It might just mean that in certain situations (in particular, when one gets good evidence against the correctness of what are in fact the correct ideals), one will end up violating some ideal or other, no matter what one ends up believing. (Christensen 2013: 91)

We could extend this line to make for a response to EB, and say that when an irrational agent enters into a disagreement she faces a genuine epistemic dilemma.²⁷ Call this the Dilemma Response.

25 Note that this expectation of Wide-Scope Conciliationism does not presuppose that the rule should tell us how to undo all possible epistemic mistakes, but only that it tell us who should conciliate.

26 Unlike with Restricted Conciliationism, it would not always be irrational to believe that we complied with Wide-Scope Conciliationism rationally. For instance, Fixit could rationally believe that she did so. But it would still be irrational for anyone who conciliates to believe that they complied with Wide-Scope Conciliationism rationally.

27 I thank Brian Weatherson for this point.

According to the Dilemma Response, when an agent *S* enters into a disagreement holding an irrational credence, a conflict between rational requirements arises.²⁸ On the one hand, *S* would be required to conciliate with her peer from that irrational credence. On the other hand, and since conciliating would not carry *S* to a rational credence (per EB), *S* would be rationally required to do something other than conciliate. So, initially irrational agents would violate some rational requirement no matter what doxastic revision they make. These agents would be in a genuine rational dilemma.

We should take the conflict in this context to be between necessary rather than sufficient conditions for believing rationally. If we take both conditions to be sufficient then satisfying either would suffice for reaching a fully rational belief, thus exposing the view to EB. Alternatively, taking one to be sufficient and the other necessary would be incoherent. If by satisfying the sufficient condition one could reach a fully rational credence, then following the other would not be necessary after all. But if both are necessary for having a rational credence then we get a genuine dilemma. No matter what *S* does, she would fall short of doing all that is rationally required of her. When no path to a fully rational credence is available to *S* then no easy path is available either, and EB is not a concern.

How bad is it to have Conciliationism imply rational dilemmas? If Christensen is right, dilemmas are already part of doxastic life, so adding a few more instances where they arise might not be so bad. What does look bad is that the Dilemma Response implies that rational dilemmas arise even on occasions where agents can apparently escape them with perfectly rational beliefs. Recall the case of Fixed Mistakes. If Fixit notices her past mistakes and corrects it by reasoning from her total evidence, she would arrive at a fully rational belief. So here too, as with Necessary Conciliationism, there is reason to doubt that conciliating is truly necessary for acquiring a rational belief, contrary to the Dilemma Response.

On the Dilemma Response, if Fixit notices her past mistake and comes to agree with her peer that *c*₂ is the justified credence, she would violate some rational requirement. Since Fixit reasons correctly from her total evidence to a belief supported by that evidence, it is hard to see what that violation would be. Perhaps one could argue that Fixit violates an alleged diachronic requirement of rational belief transition. For example, one might say that she violates the Bayesian updating rule, which sanctions as rational only belief revisions made by conditionalizing on new evidence. But the Bayesian model applies only to ideally rational agents, and does not tell irrational agents how to fix their mistakes. However, EB concerns just those cases where agents have made errors of rationality. The Bayesian model is silent on how such agents should revise, and cannot be of help here.

Moreover, while proponents of the Dilemma Response might suggest some other respect in which it is rationally criticizable to fix past errors in judgment, the opposite should strike us as true. An agent who recognizes past errors does precisely the right thing, and in precisely the right way. It is surely wrong to maintain that an agent who forms an irrational belief and then fixes it by reasoning correctly from her evidence makes two rational mistakes – first in forming the irrational belief, and then in fixing the error.

²⁸ On this account, rationality sometimes requires that we have an irrational credence.

3.3 The different ‘ought’ response

Another response to EB that is friendly to the idea that conciliation is only necessary can be derived from Miriam Schoenfield’s discussion of Calibrationism:

We should not think of calibrationism as a principle of epistemic rationality, where such principles are understood as principles about the relationship between bodies of evidence and belief states. Rather, we should think of calibrationism as a principle of reasoning. I am understanding principles of reasoning as principles about which transitions of thought, or reasoning processes, should occur in the process of deliberation. (Schoenfield 2015a: 452–3)

If we adapt Schoenfield’s idea to Conciliationism we get a view on which conciliating is required by the rules of *correct reasoning*. We can capture the point by swapping ‘rationality’ for ‘correct reasoning’ in Naïve Conciliationism:

Reasoning Conciliationism: If one has credence c_1 in p , and one encounters a peer who has a different credence c_2 in p , then correct reasoning requires one to conciliate to an intermediate credence c_3 .

On Reasoning Conciliationism, conciliating in a peer disagreement is a requirement of correct reasoning rather than a requirement of rationality. This means that in cases of disagreement conciliation is required even if the resulting credence would not be rational. Thus, although conciliation is required, a rational credence is not guaranteed for those who conciliate, and the view avoids EB.

Schoenfield’s suggestion is straightforwardly usable as a response to EB. While rationality may require us to revise our beliefs one way, correct reasoning may require us to revise another way, and no conflict between rational requirements would arise. The basic move is to distinguish between different kinds of requirements, or different senses of ‘ought,’ and say that Conciliationism derives from a normative requirement that is not a rational requirement.²⁹ Since the conciliatory requirement on this view is not a rational requirement, the view is safe from EB. The view looks preferable to the Dilemma Response, because the requirements said to conflict do not appear to give rise to a rational dilemma.

An initial concern with Reasoning Conciliationism is that the view multiplies kinds of normative requirements (or senses of ‘ought’) beyond necessity. Correct reasoning certainly seems like a kind of rational requirement, as it is often necessary for believing rationally. If an agent acquires a belief by reasoning incorrectly from the evidence, her belief would typically not be rational, even if the evidence justifies that belief. But defenders of Reasoning Conciliationism cannot grant that conciliation is necessary for rational belief, on pain of having the view reduce to Necessary Conciliationism (and deem necessary what looks like unnecessary belief revision in Fixed Mistakes).

A more serious concern is that the normative domain of reasoning requirements is obscure. If reasoning requirements are kinds of rational requirements, Reasoning Conciliationism would reduce to Naïve Conciliationism, Necessary Conciliationism, or the Dilemma Response. If reasoning requirements are not kinds of rational requirements,

29 Worsnip (2018) discusses a similar distinction between oughts of evidence-responsiveness and oughts of coherence. Distinguishing rational requirements from requirements of accuracy is another possibility here. For a related discussion see Schoenfield (2018).

we would be rationally required to ignore them when revising our beliefs. Different normative domains may require us to believe various things that are in conflict with what rationality requires. For instance, morality may require us to believe that a family member is innocent of a crime well past the point where the evidence suggests otherwise. But such a requirement does not speak to what is rational to believe, and so we are rationally required to ignore it in rational belief formation. Unless the requirement to conciliate is in service of rationality – a position that proponents of Reasoning Conciliationism must reject – it is unclear what could make ignoring it irrational.

Perhaps defenders of Reasoning Conciliationism could insist that we do away with traditional talk of rational requirements, and replace it with talk of evidential requirements and reasoning requirements. Taking evidential and reasoning requirements to represent their own respective normative domains would separate Reasoning Conciliationism from other views considered so far, thus avoiding the initial concern I mentioned. Eliminating talk of rationality would also allow the view to avoid the second concern, namely, that rationality requires us to ignore reasoning requirements if those are not rational requirements. Nevertheless, even on such a revisionary view, the two allegedly distinct kinds of requirement would not be distinct enough. Just as correctly reasoning from rational beliefs ensures that we arrive at further rational beliefs, correctly reasoning from evidentially supported beliefs ensures that we arrive at further evidentially supported beliefs. Indeed, correct reasoning seems valuable only insofar as it allows us to reach evidentially supported beliefs. Thus, whenever a conflict between evidential requirements and reasoning requirements would arise, the former would trump. If entering into a disagreement with an irrational belief creates conflict between the domains, the reasoning requirement to conciliate would lose.

Similar dialectical moves are available in response to other interpretations of Conciliationism that take it to employ a different normative operator than that of rationality. If the proposed conciliatory requirement is a kind of rational requirement, the corresponding view would reduce to one of the previously discussed responses to EB. If the proposed requirement is not a kind of rational requirement, we would be rational to ignore it. If complying with the proposed requirement is conducive to complying with rational requirements, rational requirements would win in cases of conflict.

There is one more important construal of Conciliationism that coheres with the thought that conciliation is only necessary for believing rationally. I count this construal as a distinct kind of answer to EB, since it addresses the epistemic significance of disagreement in the more general context of the rational response to our total evidence.

4. CONCILIATIONISM AND TOTAL EVIDENCE

The disagreement debate has led some to provide accounts that tell us what belief anyone who encounters disagreement should have all-things-considered.³⁰ Christensen (2016) suggests the following:

³⁰ See Sliwa and Horowitz (2015), Christensen (2016).

Idealized Thermometer Model: The credence in C it would be rational for the agent to form, given all her evidence, is the credence in C that would be rational independent of C 's support from first-order considerations, conditional on a relevantly similar agent adopting credence n in C on the basis of first-order support from the agent's evidence.³¹

Christensen mentions Kelly's EB worry after offering the Idealized Thermometer Model, and says that the model "puts rational credences" where traditional Conciliationism "had the agent's actual initial credences." So, while the Standard Thermometer Model orders agents who disagree to conciliate from their starting credences, the Idealized Thermometer Model orders them to conciliate from the credence that the pre-disagreement evidence justifies.

Rules like the Idealized Thermometer Model describe how everyone should weigh information about their *expected reliability*³² against their other evidence.³³ Since disagreement plausibly provides information of this sort, these rules include the rational response to disagreement. We may describe conciliatory rules that take this approach as sharing a general schema:

Total Evidence Conciliationism: The credence that agent S should have in p is that which S 's pre-disagreement evidence justifies, properly modified by the reliability evidence introduced by a disagreeing peer's view.

Rules that take our total evidence into account (hereafter, 'total evidence rules') cannot be accused of allowing easy bootstrapping. In order to follow them correctly an agent has to give enough weight to what the pre-disagreement evidence in fact supports, as opposed to what her own view is. As a result, these rules are far from easy to follow, and do not license two irrational agents to simply move toward one another and call it a day. Note that total evidence rules will not always recommend that agents move toward their peers. The total evidence may justify a relatively high credence even when both peers initially have irrationally low ones. In such cases, one of the agents would have to revise in a direction opposite to that of her peer's credence. But as long as a total evidence rule recommends that we give equal (or substantial) weight to the initial credences of all the relevant parties, it would be conciliatory in an important respect. In addition, total evidence rules can tell us something substantive about how evidence regarding our reliability (like disagreement) figures into rational belief.³⁴ They can do this by telling us how to weigh each kind of evidence. If so, total evidence rules may also avoid the charge of not providing insight into, or guidance about the rational response to disagreement. Total evidence rules are not simply repeating the evidentialist mantra to believe what

31 Here Christensen assumes n to be the rational credence to have given one's first-order considerations. He also assumes takes 'relevantly similar' to mean "someone who is similar with respect to the reliability evidence the agent has about herself."

32 Sliwa and Horowitz (2015) use similar terminology, and define 'expected reliability' as one's expected propensity to guess correctly.

33 More generally, these accounts tell us how to weigh our first-order evidence against our evidence regarding what that first-order evidence supports (i.e., against our *higher-order* evidence). See Christensen (2010) and Kelly (2010) for discussions of higher-order evidence.

34 See Leydon-Hardy (2016) for one such concern.

all of our evidence supports, even if they do not tell us precisely what our total evidence supports in every case.

Despite its advantages, Total Evidence Conciliationism still shares a lot with Restricted Conciliationism, including its main problems. Recall that the good news about Conciliationism came from its promise to help us hedge our doxastic bets. But including a need to correctly assess the pre-disagreement evidence anywhere in a conciliatory account again seems to get the motivation for conciliating backwards. Complying with the account would be even harder than believing what the pre-disagreement evidence alone supports. If Total Evidence Conciliationism is true, and if we struggle to rationally believe what the pre-disagreement evidence supports, we would be even less likely to rationally believe what our total evidence supports.³⁵ Here too (as with Restricted Conciliationism), this initial oddity is indicative of another issue, namely, that there is an inconsistency in following Total Evidence Conciliationism correctly and rationally believing that we did.³⁶ For suppose that we come to rationally believe that some version of Total Evidence Conciliationism is the right rule, and we also rationally believe that we followed the rule correctly. Believing that would seem to entail believing that we followed the rule's steps correctly. This includes believing that we correctly identified what the pre-disagreement evidence supports. But if we rationally believe that we correctly identified what the pre-disagreement evidence supports, we should form the corresponding credence, and not modify it in any way. In other words, we would have to roll back whatever changes we made to the credence that the pre-disagreement evidence justifies. So, believing that we followed the rule correctly is inconsistent with maintaining its recommended credence. The alternatives of disbelieving or suspending judgment about whether we followed Total Evidence Conciliationism correctly would again seem to conflict with a key motivation for conciliatory views, namely, resolving the akratic tension in having a belief and at the same time having serious doubts that the belief is rational.

Christensen defends the Idealized Thermometer Model from the charge of sanctioning epistemic akrasia:

Epistemic akrasia is not, per se, a problem at all. Thinking that a belief of yours is irrational in a particular way should disturb you – that is, give you a reason to change that belief – only insofar as the particular irrationality indicates that a different belief would have greater expected accuracy. (Christensen 2016: 416)

Applying the idea to the issue at hand, it must be that we should comply with Total Evidence Conciliationism and either disbelieve that we did – i.e., believe that the credence we end up with is irrational – or suspend judgment on the matter. In both cases we would

35 It could be that agents who try to evaluate the pre-disagreement evidence alone will tend to end up with credences that are further away from what that evidence supports than if they tried to follow Total Evidence Conciliationism. That would only show is that, in some prudential sense, we may be better off trying to follow a conciliatory strategy than not. This kind of benefit does not typically figure into rational requirements. Many things that we can do to help ourselves come closer to what our evidence supports (like staying hydrated, getting enough sleep, or triple-checking our reasoning) are not rationally required.

36 Sliwa and Horowitz (2015) raise a similar issue with their Evidential Calibration principle. Christensen (2016) applies the objection to the Idealized Thermometer Model, and offers a response I soon discuss.

have serious justified doubts that our credence is rational, but we could also maintain that credence rationally.³⁷

Had it not been irrational to hold on to a belief despite reason to think that it is irrational, we might expect there to be little pressure to change our beliefs due to disagreeing peers. On Christensen's view, though, aversion to akratic tensions may not be the source of this pressure. The pressure to change a belief due to disagreement may instead be coming from the fact that disagreement gives us reason to think that a different belief than ours has greater expected accuracy.³⁸

Be that as it may, allowing a kind of rational akrasia seems to lead to a contradiction. Grant for the sake of the argument that akrasia is not, per se, a problem, and that rationality indeed requires that we maximize something like expected accuracy. If the Idealized Thermometer Model (or some other total evidence rule) gives rise to rational instances of akrasia, we may either rationally disbelieve that we followed it correctly, or suspend judgment on the matter, despite following it correctly. For that to be true, it must be that we lack sufficient reason to believe that a different credence than the one we ended up with has greater expected accuracy – or our akratic state would not be rational. But if we know that rationality requires maximizing expected accuracy, and we should not believe that a different credence than ours does better on the accuracy front, then we should believe that we did all that rationality requires of us. This would entail believing that we followed all of the genuine rational requirements that we are under, including the Idealized Thermometer Model. So we start with the assumption that we should disbelieve that (or suspend judgment about whether) we followed the rule correctly, and conclude that we should believe that we did. Thus, granting that rules like the Idealized Thermometer Model sanction akratic doxastic attitudes lands us in a contradiction, where in some cases we both should and should not believe that we followed the rule correctly.

The thought here is similar to one already mentioned when discussing the different 'ought' response to EB. Rational beliefs and accurate beliefs cannot be in competition with one another over what we should believe. What we should believe just is what is rational to believe, even if what this amounts to is that which has greatest expected accuracy. It is therefore incoherent to believe that a belief of ours is irrational while also thinking that we should have it for the sake of accuracy. The relevant 'should' has to be that of rational belief, if it is to have any normative bite in this context. If it is not the 'should' of rational belief, it would be unclear what kind of requirement it is, and why we would not be rational to ignore it.

5. THE UPSHOT FOR CONCILIATIONISM

The Easy Bootstrapping argument shows that the pursuit of a conciliatory procedure to follow in cases of disagreement has to be sensitive to the possibility that we enter into a

37 I understand Christensen's response to amount to a defense of akrasia as *rationally* permitted in the relevant situations. Only if akrasia is not a rationality problem in those situations then we should sometimes be akratic.

38 The upcoming discussion of akrasia does not rest on some particular understanding of 'expected accuracy.' See Schoenfield (2015b) for a standard definition of 'expected accuracy.'

disagreement with an irrational credence. Since correct procedures only give good outputs when fed good inputs, agents who simply conciliate from an irrational credence will not reach an all-things-considered rational one.

It is a lot to expect a concise rule about disagreement to guide all agents in undoing their past mistakes before conciliating. This observation, however, is not grounds for dismissing the bootstrapping concern. The challenge facing conciliationists is to tell us how to conceive of the conciliatory requirement so that it does not allow bootstrapping, and is still epistemically substantive. We have seen three kinds of ways conciliationists attempt to meet this challenge: restricting Conciliationism only to rational agents, treating Conciliationism as necessary but not sufficient for fully rational beliefs, and weighing what our pre-disagreement evidence supports against the evidence introduced by the disagreement.

When we restrict Conciliationism, we are left with a bizarre rational requirement. A conciliatory view that requires only those who responded to their evidence rationally to conciliate misses its mark. When we treat Conciliationism as merely necessary for rational belief, we are left wondering how an agent should revise her beliefs when encountering disagreement. Since this is the central puzzle in the peer disagreement debate, a view that says nothing more than that is unsatisfying. Some apparently promising ways of developing the idea suggest themselves: Necessary Conciliationism, Wide Scope Conciliationism and Reasoning Conciliationism. But these suggestions face serious problems. They either reduce to Restricted Conciliationism, or deem necessary an unnecessary doxastic revision, or can be rationally ignored. A further move available to conciliationists is to bite the bullet and admit that Conciliationism sometimes leads to genuine conflicts between rational requirements. This approach has the implausible result that fixing our past mistakes is at least partly irrational.

Finally, total evidence rules may seem viable for telling us how disagreement affects rational beliefs. Still, if such rules are to avoid a bootstrapping problem, they end up facing challenges similar to those that Conciliationism faces when restricted to rational agents only. Primarily, these rules secure akratic tensions between our doxastic attitudes, and treating these tensions as rational leads to a contradiction.

It is not just Conciliationism that is vulnerable to EB. Other alleged rules about the correct response to information about our rationality are open to the charge. Consider, for example, a rule that requires agents to reduce their confidence in p by some degree upon learning that they tend to overestimate matters like p . EB would apply to such a rule just as it does to Conciliationism, and attempts to have the rule avoid the charge would send us down a dialectical path parallel to the one followed in this paper. This is some reason for thinking that we could use the arguments in this paper as part of a broader, steadfast project. Such a project would not need to involve denial of the compelling intuition that we should take seriously signs of our rational shortcomings.³⁹ But the way to take those signs seriously cannot be with a blanket revision of the credences that we happen to have.⁴⁰

39 Schechter (2013) entertains the response of double-checking our reasoning as a rationally responsible reaction to signs of our irrationality.

40 Special thanks goes to Stewart Cohen and Juan Comesaña for offering invaluable feedback and guidance through this project.

REFERENCES

- Ballantyne, N. & Coffman, E.J. 2012. 'Conciliationism and Uniqueness.' *Australasian Journal of Philosophy*, 90: 657–70.
- Broome, J. 1999. 'Normative Requirements.' *Ratio*, 12: 398–419.
- Christensen, D. 2007. 'Epistemology and Disagreement: The Good News.' *Philosophical Review*, 116: 187–217.
- 2009. 'Disagreement as Evidence: The Epistemology of Controversy.' *Philosophy Compass*, 4(5): 756–67.
- 2010. 'Higher-Order Evidence.' *Philosophy and Phenomenological Research*, 81(1): 185–215.
- 2011. 'Disagreement, Question Begging, and Epistemic Self Criticism.' *Philosophers Imprint*, 11(6): 1–22.
- 2013. 'Epistemic Modesty Defended.' In D. Christensen and J. Lackey (eds), *The Epistemology of Disagreement: New Essays*, pp. 77–97. Oxford: Oxford University Press.
- 2016. 'Disagreement, Drugs etc.: From Accuracy to Akrasia.' *Episteme*, 13(4): 397–422.
- Cohen, S. 2013. 'A Defense of the (almost) Equal Weight View.' In D. Christensen and J. Lackey (eds), *The Epistemology of Disagreement: New Essays*, pp. 98–120. Oxford: Oxford University Press.
- Elga, A. 2007. 'Reflection and Disagreement.' *Noûs*, 41(3): 478–502.
- Enoch, D. 2010. 'Not Just a Truthometer: Taking Oneself Seriously (But Not Too Seriously) in Cases of Peer Disagreement.' *Mind*, 119(476): 953–97.
- Feldman, R. 2006. 'Epistemological Puzzles About Disagreement.' In S. Hetherington (ed.), *Epistemology Futures*, pp. 216–36. Oxford: Oxford University Press.
- 2007. 'Reasonable Religious Disagreement.' In L. Antony (ed.), *Philosophers Without Gods*, pp. 194–214. Oxford: Oxford University Press.
- Finlay, S. 2010. 'What Ought Probably Means, and Why You Can't Detach it.' *Synthese*, 177: 67–89.
- Harman, G. 1986. *Change in View*. Cambridge, MA: MIT Press.
- Horowitz, S. 2014. 'Epistemic Akrasia.' *Noûs*, 48(4): 718–44.
- Huemer, M. 2011. 'The Puzzle of Metacoherence.' *Philosophy and Phenomenological Research*, 82(1): 1–21.
- Kelly, T. 2010. 'Peer Disagreement and Higher-Order Evidence.' In T. Warfield and R. Feldman (ed.), *Disagreement*, pp. 111–75. Oxford: Oxford University Press.
- 2014. 'Believers as Thermometers.' In J. Matheson and R. Vitz (eds), *The Ethics of Belief*, pp. 301–14. Oxford: Oxford University Press.
- Leydon-Hardy, L. 2016. 'Some Resistance to the Idealized Thermometer Model.' *Episteme*, 13(4): 423–6.
- Schechter, J. 2013. 'Rational Self-Doubt and the Failure of Closure.' *Philosophical Studies*, 163: 428–52.
- Schoenfeld, M. 2015a. 'A Dilemma for Calibrationism.' *Philosophy and Phenomenological Research*, 91(2): 425–55.
- 2015b. 'Bridging Rationality and Accuracy.' *Journal of Philosophy*, 112(12): 633–57.
- 2018. 'An Accuracy Based Approach to Higher Order Evidence.' *Philosophy and Phenomenological Research*, 96(3): 690–715.
- Smithies, D. 2012. 'Moore's Paradox and the Accessibility of Justification.' *Philosophy and Phenomenological Research*, 85(2): 273–300.
- Sliwa, P. & Horowitz, S. 2015. 'Respecting All the Evidence.' *Philosophical Studies*, 172(11): 2835–58.
- Titelbaum, M. 2015. 'Rationality's Fixed Point (Or: In Defense of Right Reason).' In J. Hawthorne (ed.), *Oxford Studies in Epistemology* 5, pp. 253–294. Oxford: Oxford University Press.
- Weatherston, B. Ms. 'Do Judgments Screen Evidence?'
- White, R. 2005. 'Epistemic Permissiveness.' In J. Hawthorne (ed.), *Philosophical Perspectives*, vol. 19, *Epistemology*, pp. 445–59. Malden, MA: Blackwell Publishing.
- Worsnip, A. 2018. 'The Conflict of Evidence and Coherence.' *Philosophy and Phenomenological Research*, 96(1): 3–44.

EYAL TAL is a postdoctoral fellow at the University of Cologne's center for contemporary epistemology. He earned his PhD in philosophy from the University of Arizona. His interests lie mainly in epistemology, metaethics, metaphysics, and the philosophy of science. His research focuses on epistemological issues concerning indirect evidence acquisition.
