

[This is the penultimate version. Please cite the published version.]

Conciliationism and the Peer-Undermining Problem¹

Kevin Matti Gaussein

Abstract: This paper develops a problem for conciliationism that is structurally similar to the self-undermining problem but which is immune to most of the solutions offered against it. A popular objection to conciliationism is that it undermines itself. Given the current disagreement among philosophers about conciliationism, conciliationism seems to require rejecting conciliationism. Adam Elga (2010) has influentially argued that this shows that conciliationism is an incoherent method. By recommending its own rejection, conciliationism recommends multiple, incompatible responses to the same body of experience. Many have offered solutions to this problem, including Elga himself. However, there is another undermining problem. When a peer disagrees about their own peerhood, conciliationism can require the rejection of the peer's peerhood. The result is that conciliationism recommends two incompatible responses to the same body of experience in a way that parallels the original self-undermining problem. But since many, if not all, of the proposed solutions to the original undermining problem depend on distinctive characteristics of conciliationism undermining itself, they cannot apply to a peer undermining themselves. I illustrate this by dividing the solutions that have been offered against the self-undermining problem into two categories: self-exempting solutions and non-exempting solutions. I then show how the proposed solutions under each type do not extend to the peer-undermining problem. If that is right, then conciliationism is faced with a more difficult undermining problem than has been so far acknowledged.

¹ **Acknowledgements:** I'm grateful to Jens Kipper, who encouraged me to pursue this topic. And my thanks go to Richard Feldman, Earl Conee, Jonathan Matheson, Mark Povich, David Christensen and the audience of Brown's 2022 philosophy graduate conference for discussion and helpful comments on different drafts of the paper. My gratitude also goes to Vincent Tanzil and Yusuke Satake for their many philosophical insights. Lastly, I am indebted to the late Kolja Keller, who provided extensive comments on this paper, and sharpened my thinking on this and many other areas of epistemology.

Section I: Introduction

According to *conciliationism*, when faced with peer disagreement about some proposition, you should lower your confidence that you are right. And a subject is your *epistemic peer* concerning some topic *t* when, roughly, she has the same evidence as you concerning *t* and is as good at evaluating the evidence concerning *t* as you.² A major objection against conciliationism is that it is self-undermining. Many thoughtful philosophers have argued that conciliationism is false (e.g. Kelly 2005, Titelbaum 2015, Plantinga 2000, Wedgwood 2010).³ Given such disagreement, conciliationism requires rejecting conciliationism. This is called the *self-undermining problem*.⁴ In this paper, I will argue that there is a similar but distinct problem that has been largely unaddressed. I will call this problem *the peer-undermining problem*. While the peer-undermining problem duplicates the same challenge to conciliationism, it is distinct from the self-undermining problem insofar as it is immune to many (if not all) of the solutions offered to the self-undermining problem. The upshot is that self-undermining poses a more difficult problem for conciliationism than has been so far realized.

The structure of this paper is as follows. In section two, I argue that conciliatory responses to disagreements about peers can result in peers undermining themselves. In section three, I argue that peer-undermining results in the same incoherence that self-undermining results in. In section four, I argue that the peer-undermining problem is distinct from the self-undermining problem. In section five, I divide the proposed solutions to the self-undermining problem into two categories and argue that the proposed solutions to the self-undermining problem in either category cannot

² This is following Kelly (2005, 173-175).

³ Christensen (2009) goes as far to say that “opinion on conciliationism is presently divided roughly evenly” (762).

⁴ Also called the *self-defeat problem* - e.g. Matheson (2014), Frances and Matheson (2018).

extend to the peer-undermining problem. I conclude that the peer-undermining problem is distinct from, and more challenging than, the self-undermining problem.

Section II: Conciliationism and Peer-undermining

The primary motivation for conciliationism is that it gets the right answers in clear cases. Here is a popular example from Christensen.

MENTAL MATH

Suppose that five of us go out to dinner. It's time to pay the check, so the question we're interested in is how much we each owe. We can all see the bill total clearly, we all agree to give a 20 percent tip, and we further agree to split the whole cost evenly, not worrying over who asked for imported water, or skipped desert, or drank more of the wine. I do the math in my head and become highly confident that our shares are \$43 each. Meanwhile, my friend does the math in her head and becomes highly confident that our shares are \$45 each. How should I react, upon learning of her belief?...the answer is obvious...I should lower my confidence that my share is \$43 and raise my confidence that it's \$45 (Christensen 2007, 193).

There are plenty of other motivating cases.⁵ In each case, it seems like a conciliatory response to the disagreements is rationally required. In fact, there is widespread agreement that, in cases like MENTAL MATH, you should adjust your credence in favor of your peer.⁶ Conciliationists generalize from clear cases to all cases of peer disagreement. Whenever one discovers peer disagreement, one is rationally required to conciliate in favor of one's peer.⁷

There is a range of conciliatory views. For the purposes of this paper, I will focus on the most popular version of conciliationism called the *equal weight view*, defended by Christensen

⁵ Other popular examples include Elga's (2007) imagined disagreement between two peers judging the outcome of a horse race (486), and Feldman's (2006) case of peers disagreeing about whether the dean is in the middle of the quad (207-208).

⁶ Other philosophers who agree include Matheson (2009) Pittard (2019), Bogardus (2009), Oppy (2010), and Kelly (2010). Oppy says the conciliatory response to such cases "seems inescapable" (183). Even Thomas Kelly, a prominent critic of conciliationism, acknowledges that clear cases make conciliationism "seem quite compelling" (2010, 185).

⁷ Conciliationists are frank about this. For example, Bogardus says, "I take it that concrete case intuitions like these strongly motivate the equal weight view. If it weren't for this intuitive support, arguments for the View – such as Elga's Bootstrapping Argument (486–8) – would lose much or all of their force" (2009, 326). And Fleisher says, "Conciliationism generalizes from intuitions about cases like MENTAL MATH to a general theory about the epistemically rational response to disagreement" (2021, 9914).

(2007), Elga (2007), Feldman (2006), and Bogardus (2009). The idea is that if someone is just as likely to be right about some issue as yourself, you should take their opinion with as much evidential value as your own. The idea can be applied in the following way: if you have a 0.9 credence that p and discover a peer has a 0.1 credence that p , then you should update to a 0.5 credence that p .⁸ In what follows, I will assume the equal weight version of conciliationism. But the problem I develop here occurs for other versions of conciliationism too.

The problem occurs in situations of peer disagreement about peerhood. In such cases, peers can undermine themselves. This seems to have been first noticed by Mulligan (2015):

How should I deal with a putative epistemic peer who thinks that I am not his peer but his epistemic superior? Let's imagine that I am, following Conciliationism, factoring in the opinion of Jimmy, whom I believe to be my epistemic peer on some matter p about which we disagree...Imagine that Jimmy says this: "Tom, I know that I disagree with you about all kinds of things, including p , but I'm not your intellectual equal. You have good reason to ignore my opinions and stick to your guns." (75)

Let's refer to Mulligan in this case as *Tom*. And let's further say that Tom justifiably believes that Jimmy is Tom's peer concerning their intellectual abilities. What should Tom believe given the peer disagreement with Jimmy about Jimmy's peerhood? According to conciliationism, Tom is rationally required to decrease his credence that Jimmy is his peer. Let's say that Tom started with a credence of 0.9 that Jimmy is his peer and that Jimmy has a credence of 0.1 that he is Tom's peer. Thus, conciliationism requires Tom to adjust his credence to 0.5. Let's say that Tom conciliates. The result is that Tom no longer believes that Jimmy is his peer.⁹ Tom now seems to be in an odd situation. His new belief state was conditioned on Jimmy's peer disagreement but no longer includes believing that Jimmy is his peer.

⁸ Christensen (2007) argues for this sort of application of the equal weight and calls it "splitting the difference" (203-204).

⁹ This is assuming that a degree of confidence of .5 is not high enough to count as belief.

Here's one reason this is an odd result. By conciliating, it seems like Tom now rejects a piece of evidence used in conciliating. Before conciliating, Tom's evidence for conciliating included his belief that Jimmy is his peer. After conciliating, Tom no longer believes that Jimmy is his peer. As a result, it seems like Tom's updated belief lacks evidential support. Thus, in peer-undermining cases, applying conciliationism can result in one's evidence no longer supporting one's updated belief state. And insofar as one's justification depends on one's evidence, one's updated belief state is no longer justified.

Mulligan (2015) says that this results in a paradox (75). But he suspects that a proper analysis of peerhood will solve the paradox. He says, “[conciliationism], which is otherwise compelling, needs to provide a conceptual analysis of epistemic peerhood that resolves these apparent paradoxes to our satisfaction” (77). However, it is hard to see how an improved analysis of peerhood can help. There are already a number of distinct notions of peerhood in the literature.¹⁰ And none of them seems capable of resolving the problem. Furthermore, a direct disagreement about peerhood is not necessary to generate the problem. So long as the disagreement has certain rational implications for one's belief about who is their peer, then the puzzle still occurs. For example, imagine a peer disagreement, not about whether someone is in fact one's peer, but about an epistemic point concerning our ability to recognize peerhood. Say you believe that your friend is your peer about the topic of peerhood. And then you discover that they believe that we cannot rationally believe that anyone is our peer. They might say, “Whether someone has just as much evidence as you and is just as competent at evaluating evidence as you about some topic is really complicated. Given our cognitive limitations, peerhood is inscrutable. That is to say, we cannot

¹⁰ For example, Kelly develops a concept of an epistemic peer in his (2005, 173-175). Feldman (2007, 144) and Christensen (2007, 188-189) give similar definitions. And Elga (2007, 484 and 499 n. 21) gives what he calls a “nonstandard” use of an epistemic peer. King (2012) discusses some of the conditions required for two subjects to count as peers.

rationally believe that someone is our peer.” After updating, it seems like you are rationally required to suspend your belief that your friend is your peer. If so, then the peer-undermining problem occurs. And it seems like that can happen no matter what analysis of peerhood one offers.

Section III. Peer-undermining Results in Incoherence

In the last section, I illustrated how conciliationism can require the rejection of a peer’s peerhood. That is, in certain cases of peer disagreement about peerhood, conciliationism recommends rejecting a peer’s peerhood. So far, all we’ve seen is that there is a sort of oddity about the updated belief state. But it isn’t clear how problematic this is. After all, what else are subjects like Tom supposed to believe?¹¹ However, there is a deeper problem in the neighborhood. I will argue in this section that the real significance of peer-undermining is that it duplicates what is perhaps the most challenging problem for conciliationism.

Recall back to the self-undermining problem. The issue is supposed to be that, given widespread disagreement about conciliationism, conciliationism seems to require the rejection of conciliationism. One might wonder what the problem here is supposed to be. That conciliationism cannot be consistently believed in certain conditions does not entail that it is false (Matheson 2014). Elga (2010) argues that the real problem with self-undermining is that it shows that conciliationism is incoherent. According to Elga, conciliationism is part of one’s inductive method.

¹¹ A similar thought is expressed in Knoks’s (2022) response to Mulligan. Knoks develops a framework of defeasible logic that aims to make the implications of conciliationism clearer. But Knoks notices that the framework does not straightforwardly solve Mulligan’s case of a peer who undermines their own peerhood. Instead, Knoks argues that such a case is not unique to conciliationism (19). For example, it occurs whenever someone testifies that p where p implies that the person is unreliable. Further, he argues that a more general framework can account for situations where one should conciliate despite no longer believing that the person is one’s peer (21). Knoks’s points seem reasonable against Mulligan’s formulation of the problem. But as I will argue in the rest of this section, a deeper problem has to do with conciliationism (together with other plausible principles of rationality) prescribing conflicting recommendations in situations that combine disagreements about peerhood with other disagreements.

An inductive method is “one’s fundamental method for taking evidence into account” (2007, 6).¹² Two methods are competitors if they offer contrary recommendations to the same course of experience. Since conciliatory views can recommend their own rejection, it follows that they are able to recommend another method. If that other method is a competing one, then there are situations where an inductive method that includes conciliationism can recommend incompatible responses. And, as Elga puts it, “it is incoherent for an inductive method to recommend two incompatible responses to a single course of experience” (2010, 7).

Elga illustrates this incoherence with consumer magazines. Imagine that Consumer Reports recommends buying only toaster X and Smart Shopper recommends buying only toaster Y. Furthermore, imagine that Consumer Reports recommends following only the advice of Smart Shopper. What you get is this:

- Consumer Reports says: “Buy only toaster X.”
- Smart Shopper says: “Buy only toaster Y.”
- Consumer Reports says: “Follow the advice of Smart Shopper.”

Consumer Reports is making recommendations that are jointly impossible to follow. It recommends buying only toaster X and only toaster Y (through recommending Smart Shopper). It is impossible for one to consistently follow the advice of Consumer Reports. Since it is recommending two incompatible responses, Consumer Reports is incoherent.

In the same way, a conciliatory view that calls for its own rejection is incoherent. Let’s say that beliefs X and Y are incompatible in the sense that they cannot both be rationally believed. Now imagine an inductive method M that recommends “in response to experience E, adopt belief X.” Imagine another method N that recommends “in response to E, adopt Y.” Now suppose that M also recommends adopting inductive method N over M. What you get is this.

¹² Elga defines a fundamental method as “one whose application is not governed or evaluated by any other method” (2007, 11, note 11).

- M: adopt belief X.
- N: adopt belief Y.
- M: follow inductive method N instead of M.

The idea here is that conciliationism recommends X and, in self-undermining cases, recommends a method that recommends belief Y.

So far, this does not result in incompatible recommendations. M recommends X and following N. It does not also recommend Y. But assumed in Elga's discussion is a plausible principle sometimes referred to as "level-connection."¹³ According to level-connection, one's credences are constrained by what credences one believes one is rational to have. So, roughly, if one believes that one should not believe p, then one should not believe p. So long as M includes a level-connecting principle, then it will also recommend adopting belief Y. Thus, assuming level-connection, M recommends X and Y. But X and Y are incompatible. So it is impossible to follow both recommendations. M is therefore incoherent.

One might be tempted to avoid this incoherence from sticking to conciliationism by rejecting level-connection. But Christensen (2013) points out that a conciliationist cannot simply reject level-connection out of hand.¹⁴ This is because level-connection seems to be part of the motivation for conciliationism. Peer disagreement about p provides evidence that one has made an epistemic mistake about p. Given such evidence, level-connection rationally requires one to adjust their credence about p. Otherwise, it wouldn't seem like conciliating is the right response in clear cases like MENTAL MATH.

¹³ This is also called "level-bridging." Christensen (2013) and Littlejohn (2014; 2020) point out that Elga's argument assumes some sort of level-connection principle. See Whiting (2020) for an overview of recent work on level-connection and Horowitz (2014) for arguments in favor of level-connection.

¹⁴ Though as we will see in section 5, Christensen has recently argued that there are situations where level-connection is not rationally required.

This problem, argues Elga, “arises in any case in which an inductive method recommends a competing method over itself” (2007, 8). He considers this to be a decisive objection against conciliationism – “Bottom line: The self-undermining problem shows that conciliatory views on disagreement should be rejected” (2010, 8).¹⁵

This point about incoherence is perhaps the most discussed objection to conciliationism.¹⁶ There have been a slew of solutions to self-undermining offered in response. However, the same incoherence that results from self-undermining extends to peer-undermining without the proposed solutions so far offered to the first extending to the second. In other words, the same incoherence occurs because peer-undermining can result in incompatible responses just as much as self-undermining. But, as I will argue later, the solutions so far offered cannot work against peer-undermining because they all rely on distinctive features of self-undermining. I will discuss those features in the next section. For now, I will argue that peer-undermining also results in conciliationism recommending incompatible responses to the same course of experience.

Here’s the idea. Imagine that your peer recommends X. Now imagine that your peer also recommends believing that they are not a peer. What you get is this:

- Peer: adopt belief X.
- No-peer: adopt belief Y.
- Peer: believe No-peer.

An inductive method that includes conciliationism will recommend belief X and recommend following No-peer. Given level-connection, one’s inductive method will also recommend belief Y. But X and Y are incompatible. Thus, one’s inductive method results in incompatible recommendations.

¹⁵ As we will see later, his solution is to offer a “partly conciliatory view” that excludes conciliationism from applying to itself.

¹⁶ Christensen (2013) calls Elga’s problem “a very powerful challenge” to conciliationism (84).

We can illustrate this by imagining a typical case of peer disagreement that occurs in conjunction with peer-undermining. And we can do that by supplementing Mulligan's case of peer-undermining in the following way:

DISAGREEING PEERS

Tom and Jimmy have spent the summer studying the debate on epistemic internalism and externalism. Tom and Jimmy are peers about the debate and about their intellectual abilities. At the end of the summer, Tom concludes that internalism is false. But, to his surprise, he discovers that Jimmy came to the opposite conclusion and believes that internalism is true. Tom conciliates and withholds his belief that internalism is false. The following day, Jimmy says to Tom, "Tom, I know that I disagree with you about all kinds of things, including internalism, but I'm not your intellectual equal. You have good reason to ignore my opinions and stick to your guns."

We can further stipulate that Tom justifiably believes that Jimmy is his peer both about the debate and about their intellectual abilities. Thus, conciliationism recommends that Tom rejects his belief that internalism is false on the basis of his peer disagreement with Jimmy. Then it also recommends rejecting his belief that Jimmy is his peer. Given that he no longer believes that Jimmy is his peer, Tom's inductive method seems to require believing that internalism is false again. After all, Tom rejected it to begin with because he considered Jimmy to be a peer. Thus, depending on whether Jimmy is a peer, Tom's inductive method recommends:

- Peer Jimmy: do not believe that internalism is false.
- No-peer Jimmy: believe that internalism is false.
- Peer Jimmy: follow No-peer Jimmy.

What an inductive method that includes conciliationism recommends if Jimmy is a peer is incompatible with what it recommends if Jimmy is not. Tom cannot consistently follow the recommendations. The problem is that an inductive method that includes conciliationism recommends following No-peer Jimmy because it recommends following Jimmy. Thus, an inductive method that includes conciliationism will recommend following Jimmy and thereby recommend incompatible responses to the same course of experience.

It follows that an inductive method that includes conciliationism is incoherent. But the reason it is incoherent here is not because conciliationism recommends against following conciliationism, but rather because it recommends following a peer who recommends against following herself. By recommending following a peer's recommendation against herself, one's inductive method results in recommending incompatible beliefs to the same course of experience. Therefore, if self-undermining results in incoherence, then so does peer-undermining.

Let me anticipate one natural objection at this point. One might think that conciliationism only applies when one believes that someone is a peer. Thus, in *DISAGREEING PEERS*, conciliationism no longer recommends adopting a higher credence in internalism after Jimmy updates on Tom's disagreement about Tom's peerhood. But that is false. Just as conciliationism does not apply only to those who have beliefs about conciliationism, so it does not apply only to those who have beliefs about peerhood. Instead, what is important is that one has (something like) evidence or justification to believe that a subject is one's peer and that they disagree. So long as one has evidence of peer disagreement, then conciliationism requires conciliating. That is consistent both with believing that conciliationism is false and believing that the subject is not one's peer. It's not as if you can escape the commitments of conciliationism just by, for every case of peer disagreement, deciding to believe that the person is not your peer.¹⁷ As I will describe in the next section, this shows that in both the self-undermining and peer-undermining problem, the disagreements about conciliationism and peerhood are actually misleading. But misleading

¹⁷ Furthermore, while this paper is focused on peer disagreement, conciliationists (e.g. Elga 2007; Mulligan 2021) have also pointed out that the rational implications of conciliationism apply not just when one discovers disagreement with one's epistemic peers, but also when one discovers disagreement with one's intellectual superiors and inferiors. The requirements of conciliationism is, thus, gradable among a spectrum of cases – it requires adjusting one's confidence more so when disagreeing with an epistemic superior, and less so when disagreeing with an epistemic inferior. My thanks to a referee for bringing this point to my attention.

evidence is still evidence. So long as one can be wrong about what their evidence supports, then it seems like the problem gets off the ground.

Section IV: Two Higher-order Problems

At first glance, the peer-undermining problem is similar to the self-undermining problem. In both cases, something is undermined by conciliating on peer disagreement. In one - conciliationism undermines itself, in the other - the peer undermines herself. And both cases of undermining seem to result in a sort of *higher-order* problem. There is a general consensus that peer disagreement involves higher-order evidence (Feldman 2005; 2009; Kelly 2005; 2010; Christensen 2010; Matheson 2009; Whiting 2020). Recall that a peer is someone who, roughly, has the same evidence as you and is just as good at evaluating the evidence as you. If you come to conclusion C on the basis of evidence E, and your peer comes to the conclusion not-C on the basis of E, then your peer's disagreement is evidence that you have misevaluated E. Thus, discovering peer disagreement gives you evidence that you made a mistake about what your evidence supports.¹⁸

This sort of evidence is “higher-order” because it is evidence that you have misevaluated the evidence.¹⁹ So in cases where one discovers peer disagreement about conciliationism, one gets evidence (i.e. higher-order evidence) that one has misevaluated the evidence about conciliationism. In cases where one discovers peer disagreement about peerhood, one gets evidence (i.e. higher-order evidence) that one has misevaluated the evidence about peerhood. And in both situations, the problem occurs when one's higher-order evidence is misleading in the sense that one's higher-order evidence supports believing a falsehood. In the self-undermining problem, we assume that conciliationism is true and then consider a situation where conciliationism requires

¹⁸ There are contested ways of understanding how this higher order evidence affects your first-order evidence. For a summary of recent developments on this issue, see (Whiting 2020).

¹⁹ Though it also counts as first-order evidence in the normal testimonial way.

that one rejects conciliationism. And in the peer-undermining problem, we assume that a subject is one's peer and then consider a situation where conciliationism requires that one rejects that person's peerhood. But believing that conciliationism is false does not make conciliationism false. There is still a fact about what conciliationism recommends given peer disagreement. In the same way, believing that it is false that a person is one's peer does not make it false that that person is your peer. There is still a fact about what conciliationism recommends given peer disagreement. Thus, both situations are higher-order problems in that what one is rationally required to do conflicts with what one rationally believes one ought to do.

Despite these similarities, the two problems are distinct in important ways. Self-undermining involves getting higher-order evidence against how one takes higher-order evidence into account. So it is a unique type of disagreement and differs from typical cases of getting higher-order evidence against one's first-order evidence. In contrast, the peer-undermining problem is just a typical instance of getting higher-order evidence against one's first-order evidence.²⁰ Further, whereas conciliationism is a rule of rationality and is (in some respect) non-empirical, peerhood is a contingent, empirical matter. So self-undermining involves disagreement about a rule of rationality and peer-undermining does not.

While the peer-undermining problem has been mostly unaddressed in the literature, the self-undermining problem has not. Given the similarities between the problems, it is tempting to think that some solution offered against the latter can solve the former. But the above differences are so significant that many (if not all) of the most plausible solutions on offer to the self-undermining problem fail to extend to the peer-undermining problem.

²⁰ Mulligan also makes this point - "Since these propositions [about peerhood] ... are no different (in any important way) from the archetypical propositions in the disagreement debate, Conciliationism faces a significant theoretical challenge" (2015, 68).

Section V: Proposed Solutions to Undermining

Proposed solutions to Elga’s self-undermining problem include Elga (2010), Christensen (2013), Pittard (2015), Matheson (2014), Bogardus (2009), Christensen (2021), Littlejohn (2020), and Knoks (2021). But many of these proposed solutions, if not all, utilize distinctive characteristics of the self-undermining problem that do not characterize disagreements about peerhood. So they make use of resources unique to the self-undermining problem that are unavailable for solving the peer-undermining problem. In order to show this, it will be helpful to classify these proposed solutions into distinct categories.

We can categorize them by how they respond to disagreement about conciliationism. Recall the incompatible recommendations that Elga says an inductive method that includes conciliationism results in:

- (a) M: adopt belief X.
- (b) N: adopt belief Y.
- (c) M: adopt inductive method N instead of M.

One cannot consistently follow both M and N. So a solution needs to show that one of the incompatible beliefs (X or Y) does not get recommended. The strategy, then, is to exclude (a) or (b). Given level-connection, the inclusion of (c) requires the inclusion of (b). So one strategy is to exclude (c) by arguing that conciliationism is exempt from conciliating on disagreements about itself (Bogardus 2009; Elga 2010; Pittard 2015). Call such strategies *self-exempting solutions*.²¹ Strategies that allow conciliationism to reject itself can be called *non-exempting solutions*.²² Non-exempting solutions aim to exclude (a) or (b) depending on whether they accept level-connection (Matheson 2014; Christensen 2013) or not (Christensen 2021; Littlejohn 2020; Knoks 2021).

²¹ The description of “self-exempting” comes from Blessenohl (2015) and Pittard (2015).

²² By describing these strategies as “solutions,” I do not mean to suggest that I think any of them successfully solves the self-undermining problem. These strategies should be understood as “proposed solutions.”

In what follows, I will consider one proposed solution for each option and argue that they fail to extend to the peer-undermining problem. And the arguments I offer that these solutions do not extend to the peer-undermining problem generalize to other proposed solutions. The point here is not that any of these solutions succeed or fail against the self-undermining problem. Neither is it that there is no solution to the peer-undermining problem. Rather, the point is that even if these solutions solve the self-undermining problem, they do not extend to the peer-undermining problem. The upshot is that the peer-undermining problem is a more challenging problem for conciliationism than the self-undermining problem.

5.1 Self-exempting Solutions

According to self-exempting solutions, conciliationism is exempt from conciliating on disagreements about conciliationism. The strategy is to eliminate (c) from the set, and as a result the incompatible belief in (b) does not get prescribed. Thus, the goal is to exclude (c):

- (a) M: adopt belief X.
- (b) N: adopt belief Y.
- ~~(c) M: adopt inductive method N instead of M.~~

In this case, M and N prescribe incompatible beliefs, but M does not prescribe following N. A number of strategies have been offered to explain why conciliationism can be self-exempt, but they all trade on a unique feature of the self-undermining problem – namely, that it is the only case where the principle requiring conciliating is what is under dispute. And, plausibly, this is true of all self-exempting solutions. After all, a solution to the self-undermining problem should not exempt just any disagreement, otherwise it risks losing the motivation for conciliationism to begin with. So there needs to be a principled reason why only conciliationism is exempt. And it seems like the only reason is that disagreeing about conciliationism is a unique type of disagreement. But, as I will argue, the peer-undermining problem does not share this feature. Consequently, such

strategies cannot extend to disagreements about peers. I will illustrate this by considering Elga's (2010) own response to the problem.

Elga's strategy is to reject the possibility that conciliationism can reject itself because it is part of one's fundamental method for taking evidence into account. And "it is in the nature of giving consistent advice that one's advice be dogmatic with respect to its own correctness" (11). Such dogmatism is a general constraint on all fundamental methods. "This general constraint provides independent motivation for a view on disagreement to treat disagreement about disagreement in a special way" (11). Otherwise, one's inductive method will result in recommending incompatible responses to the same body of evidence, which is incoherent. Thus, Elga endorses what he calls a "partially conciliatory" view (Elga 2010, 8). When one is confronted with disagreement about anything other than disagreement, conciliation is required. But when one is confronted with disagreement about disagreement, conciliation is not required.

However, this move cannot help the peer-undermining problem. It is consistent with the peer-undermining problem that a conciliationist is dogmatic about conciliationism. After all, the peer-undermining problem does not involve gaining evidence against how one treats evidence, nor does it involve gaining evidence against any of one's fundamental methods. In this sense, there is nothing "special" about the peer-undermining problem. It is just another case of getting higher-order evidence against one's first-order evidence. Even if there is reason to be dogmatic about conciliationism, there is no reason to be dogmatic about peerhood.

One might hope to apply Elga's point by making it more general in the following way. In order for one's fundamental methods to work, one cannot allow the potential for any undermining - including from peers. Perhaps it is in the nature of giving consistent advice that one's fundamental method should be dogmatic about anything that can undermine itself.

The problem is that generalizing Elga's strategy to apply to any undermining goes too far. Peerhood is an empirical matter that one can get more or less evidence about. Your friend might be just as good as you at evaluating evidence while sober, and not while drunk. It is just not the sort of thing that one can be dogmatic about. Further, being dogmatic about peerhood means that one cannot conciliate in non-undermining cases of disagreeing about peers. For example, say that I justifiably believe Jim is not my peer but Zach is. And let's say I find out that Zach believes that Jim is my peer. Intuitively, I should increase my credence that Jim is my peer to some degree. And why not? To not do so goes against the motivation for conciliating in clear cases of peer disagreement like MENTAL MATH. Thus, it does not seem profitable to extend Elga's strategy to anything capable of self-undermining.

5.2 Non-exempting Solutions

Recall again the set of incompatible recommendations that Elga says an inductive method that includes conciliationism results in:

- (a) M: adopt belief X.
- (b) N: adopt belief Y.
- (c) M: adopt inductive method N instead of M.

Unlike self-exempting solutions, non-exempting solutions allow conciliationism to result in its own rejection. That is to say, non-exempting solutions accept (c). In this case, M can prescribe following N. Given level-connection, if one believes that one should believe Y, then one should believe Y. Thus, non-exempting solutions can be divided by whether they accept level-connection or not. I will consider both options.

5.2.1 Non-exempting with Level-Connection

Since they accept (c), non-exempting solutions that accept level-connection are committed to (b). So the strategy is to exclude (a).

- (a) ~~M: adopt belief X.~~
- (b) N: adopt belief Y.
- (c) M: adopt inductive method N instead of M.

In this case, M can prescribe following N. But by doing so, M no longer prescribes belief X. And so conciliationism does not result in incompatible recommendations to the same course of experience. Here too, non-exempting solutions to the self-undermining problem make use of distinct aspects of the self-undermining problem. As a result, they do not apply to the peer-undermining problem. I will illustrate this by considering Matheson's (2014) proposed solution.²³

Matheson (2014) argues that it is not incoherent for conciliationism, in certain contingent conditions, to recommend rejecting itself. The incoherence is supposed to occur because there are multiple prescriptions to the same body of evidence. There are multiple competing prescriptions as a result of there being competing rules about how to respond to the evidence. Matheson argues that if there were a rule about which prescription to follow (a meta-rule), then the problem would be solved. And he thinks there is such a rule. The idea is that prescriptions about rules have a higher priority over other prescriptions. Thus, in cases where conciliationism requires its own rejection, the prescription to reject conciliationism is a "higher-order" prescription.²⁴ And Matheson argues that this makes it "upstream" to other prescriptions (7). Like a raft that floats down a branching river, what happens downstream depends on which branch of the river the raft floats down. In the same way, M's prescription to "adopt inductive method N instead of M" is upstream to the prescriptions to adopt belief X and Y. Thus, conciliationism only prescribes belief

²³ Another solution that plausibly fits under this category is Christensen's (2013) conflicting ideal approach. Though, as we will see, Christensen no longer defends this solution.

²⁴ Important to note that this notion of higher-order prescriptions is distinct from the notion of higher-order evidence.

Y, not both X and Y.²⁵

The problem is that this does not extend to peer-undermining cases. In the self-undermining problem, the disagreement is about incompatible rules about how to respond to a certain kind of evidence. Thus, a meta-rule, or a rule about which rule to follow, might be able to resolve conflicting prescriptions. But there are no conflicting rules in the peer-undermining case. The peer-undermining problem does not involve a disagreement about how to take evidence into account. Rather, it is a disagreement about peerhood. Whereas the self-undermining problem results in a rule being disputed, the peer-undermining problem results in the application of the rule being disputed. And since that does not involve conflicting rules, a meta-rule cannot help.

Here's another way to put it. We can distinguish between some rule R and some circumstance C in which R is applicable in the following way. R says that when in C, believe proposition p. Using this lingo, we can state the self-undermining problem as a conflict between two rules:

- R1: when in C, believe p.
- R2: when in C, believe ~p.

Here a meta-rule for choosing between R1 and R2 might help. But rather than conflicting rules being at odds, the conflict in peer-undermining cases comes from which circumstance one is in:

- R: when in C1, believe p.
- R: when in C2, believe ~p.

Here, the same rule remains constant – specifically, when faced with peer disagreement, conciliate in favor of the peer. The dispute is whether some disagreement counts as a peer disagreement. And so long as one can be rationally mistaken about which circumstances one is in, then the problem

²⁵ Matheson's strategy is undergirded by evidentialism. Littlejohn (2014) argues that Matheson's evidentialism combined with level-connection results in the same sort of incompatible requirements that Egla argues conciliationism results in. Matheson (2015) responds to this charge.

gets off the ground.

To illustrate, we can think of C1 as being the type of situation where one is faced with peer disagreement. In DISAGREEING PEERS, Tom conciliates and increases his confidence in internalism. Tom's disagreement with Jimmy is analogous to Tom's being in C1. In the face of the second disagreement with Jimmy, conciliationism requires Tom to no longer believe that Jimmy is his peer. This is analogous to Tom believing that he is in C2. Thus, Tom is actually in C1, and so conciliationism recommends *p*. But conciliationism requires Tom to believe that he is in C2. The result is that conciliationism makes incompatible recommendations. And unlike self-undermining cases, there is only one rule at play here. So a rule about which rule to follow will not help.

5.2.2 Non-exempting without Level-Connection

Now consider non-exempting solutions that do not accept level-connection. Given that they allow conciliationism to reject itself, they can include (c). But unlike other non-exempting solutions, these allow the violation of level-connection. That is, even if one believes that one should believe *Y*, it can still be rational for one to not believe *Y*. Thus, this strategy can aim to exclude (b) instead of (a):

- (a) M: adopt belief *X*.
- ~~(b) N: adopt belief *Y*.~~
- (c) M: adopt inductive method N instead of M.

Such a strategy has been unpopular because, as noted in section 3, conciliationism seems to be partly motivated by level-connection. So it seems odd to use a strategy in defense of conciliationism that requires rejecting a principle that motivates conciliationism. Nonetheless, such a strategy has recently gained some traction (e.g. Christensen 2021; 2022; Littlejohn 2020). The idea, then, is that level-connection applies in most cases but not all. And when level-

connection does not apply, conciliating is not required. Such cases include disagreements about conciliationism.

Similar to self-undermining solutions, it seems like non-exempting solutions that reject level-connection must trade on distinctive features of the self-undermining problem. If they don't, then they run the risk of undercutting the motivation for conciliationism to begin with. That is to say, there needs to be a reason why level-connection is not required after updating on disagreements about conciliationism but is required in other cases. However, since they trade on features of the self-undermining problem that the peer-undermining problem does not share, this version of the non-exempting solution cannot extend to the peer-undermining problem. I will illustrate this with Christensen's recent response to the self-undermining problem.

Christensen (2021) offers what can be called an *akratic* solution to the self-undermining problem. Epistemic akrasia can be characterized as believing two propositions – p and *one should not believe p* . The two propositions conflict by level-connection. So one can only affirm the rationality of epistemic akrasia if one rejects level-connection. Christensen argues that epistemic akrasia is irrational in most situations. But in certain situations where what one believes is rational comes apart from what one believes is accurate, epistemic akrasia can be rational. Christensen illustrates this with a college student named Dara (2021, 2198-2203).²⁶ Dara's professors all argue for a (false) theory of rationality called Deductive Purism according to which induction provides no reason to think a belief is rational, but can still provide reason to think it is accurate. Given the expertise of Dara's professors, Dara might be rationally required to believe Deductive Purism. Then when considering whether his sandwich will nourish or poison him, Dara realizes that he only has inductive reasons for thinking it will nourish rather than poison him. It seems like Dara

²⁶ Christensen offers other examples of intuitively rational epistemic akrasia in (2021) and (2022).

is in a situation where he is rational to believe that i) his sandwich will nourish him and ii) that it is not rational to believe that his sandwich will nourish him. He is rational to believe (i) because he has good rational grounds to believe (i). After all, Deductive Purism is false. And further, he is rational to believe (ii) because he rationally believes Deductive Purism.

Christensen says that something similar occurs for cases of disagreements about conciliationism. Conciliationism is a rule of rationality. So when one conciliates on peer disagreement about conciliationism, one no longer believes that credences conforming to conciliationism (say, about internalism) are rational. But one may, nonetheless, continue holding such credences because one rationally believes that the credences are accurate. After all, they were updated on the disagreements of someone whose judgments are just as accurate. Thus, according to Christensen, after conciliating on peer disagreement about conciliationism, it can be rational both to have a .5 credence in internalism and believe that such a credence is irrational.

The problem is that the above situation does not occur in cases of disagreements about peerhood. Someone is your peer when they are, roughly, just as good as you are at evaluating evidence and have the same evidence you do. That is, a peer is someone who is just as reliable as you are about some topic. Consider DISAGREEING PEERS again. After updating on Jimmy's disagreement about his peerhood, Tom should now think his credence about internalism is irrational. But it is not because he changed his mind about what counts as rational. Instead, it is because he thinks his credence was formed on an unreliable basis. He updated it on what he believed was peer disagreement (as it turns out, it really was peer disagreement). But, after conciliating on Jimmy's disagreement about peerhood, he rationally believes that it was not a peer disagreement. So what he believes is rational does not come apart from what he believes is accurate.

Here's another way of making the point. Christensen (2021; 2022) argues that in cases of rational epistemic akrasia, the subject should be able to reason about each proposition in a way that makes the divergence of accuracy and rationality explicit. In the case of Dara, Dara can reason about the second conjunct like this: "I should believe that the belief that my sandwich will nourish me is irrational because there is no non-circular way to justify induction."²⁷ But Dara can still reason about the first conjunct in the following way: "I should believe that my sandwich will nourish me because induction can support accurate beliefs, even if such beliefs are irrational." And similarly, the subject that discovers peer disagreement about internalism and conciliationism can reason, "I should believe that a 0.5 credence in internalism is irrational because I no longer believe in conciliationism. But I should maintain a 0.5 credence in internalism because the person who disagreed with me is just as reliable as I am about internalism, even if such a credence is not the rational one to have." The problem is that Tom cannot rationally reason in an analogous way. Tom's reasoning about the second conjunct should be something like this: "I should believe that my credence in internalism is irrational because it was conditioned on Jimmy being just as reliable as I am, which I no longer believe." And Tom cannot appeal to accuracy to support maintaining his credence in internalism. Instead, he should reason, "Since I no longer think Tom is just as reliable as I am, my credence in internalism is inaccurate." Whereas conciliationism is about rationality, peerhood is about accuracy. More precisely, peerhood has to do with whether a subject is just as good at evaluating the evidence as you are (i.e. whether they are just as *accurate* as you are). So by rejecting Jimmy's peerhood, Tom cannot appeal to something like accuracy in his reasoning to maintain his credence in internalism. Thus, even if accuracy and rationality come apart in cases of disagreements about conciliationism, they do not in cases of disagreements about

²⁷ This is how Dara's professor tells Dara to reason (Christensen 2021, 2199).

peerhood.

Conclusion

I conclude that the peer-undermining problem is a distinct problem from, and more challenging than, the self-undermining problem. Proposed solutions to the self-undermining problem can be divided into two categories: self-exempting and non-exempting solutions. There seems to be a principled reason that self-exempting solutions always depend on distinctive features of the self-undermining problem – namely, they need to explain why only disagreements about conciliationism are exempt from conciliating. So they cannot extend to the peer-undermining problem because disagreements about peerhood are just ordinary cases of disagreement. In contrast to self-exempting solutions, some non-exempting solutions need not rely on the unique features of disagreements about conciliationism. Thus, they have a better chance of extending across both problems. However, insofar as they do rely on unique features of disagreeing about conciliationism, they too cannot extend. Moreover, those non-exempting solutions that reject level-connection, like self-exempting solutions, also seem to have a principled reason to always depend on distinctive features of the self-undermining problem. In particular, they need to explain why level-connection applies to disagreements in general but not to belief states conditioned on disagreements about conciliationism. Therefore, the peer-undermining problem shows that an inductive method that includes conciliationism is incoherent in the same way that the self-undermining problem does, but without allowing many (if not all) solutions so far offered against the latter to apply to the former.

Works Cited:

- Blessenohl, Simon. (2015). Self-Exempting Conciliationism is Arbitrary. *Kriterion: Journal of Philosophy*, 29(3), 1-22.
- Bogardus, Thomas. (2009). A Vindication of the Equal-Weight View. *Episteme*, 6(3), 324-335.
- Christensen, David. (2007). Epistemology of Disagreement: The Good News. *Philosophical Review*, 116(2), 187-217. doi:10.1215/00318108-2006-035
- Christensen, David. (2009). Disagreement as Evidence: The Epistemology of Controversy. *Philosophy Compass*, 4(5), 756-767.
- Christensen, David. (2010). Higher-order Evidence. *Philosophy and Phenomenological Research*, 81(1), 185-215.
- Christensen, David. (2011). Disagreement, Question-Begging and Epistemic Self-Criticism. *Philosophers Imprint*, 11(6): 1–22.
- Christensen, David. (2013). Epistemic Modesty Defended. In D. Christensen & J. Lackey (Eds.), *Epistemology of disagreement: New Essays*. Oxford: Oxford University Press. doi:10.1093/acprof:oso/9780199698370.001.0001
- Christensen, David. (2019): Formulating Independence. In *Higher-Order Evidence: New Essays*, M. Skipper & A. Steglich-Petersen (Eds.), Oxford: Oxford University Press.
- Christensen, David. (2021). Akratic (Epistemic) Modesty. *Philosophical Studies*, 178(7), 2191-2214.
- Christensen, David. (2022). Epistemic Akrasia: No Apology Required. *Noûs*. <https://doi.org/10.1111/nous.12441>
- Elga, Adam. (2007). Reflection and Disagreement. *Noûs*, 41(3), 478-502.
- Elga, Adam. (2010). How to Disagree about how to Disagree. In R. Feldman and T. Warfield (Eds.), *Disagreement*. Oxford: Oxford University Press.
- Feldman, Richard. (2005). Respecting the Evidence. *Philosophical Perspectives*, 19(1), 95-119. doi:10.1111/j.1520-8583.2005.00055.x
- Feldman, Richard. (2006). “Epistemological Puzzles about Disagreement.” In Stephen Hetherington (ed.), *Epistemology Futures*. Oxford: Oxford University Press, pp. 216–36.
- Feldman, Richard. (2007). “Reasonable religious disagreements.” In Louise Antony (ed.), *Philosophers Without Gods: Meditations on Atheism and the Secular Life*. Oxford University Press. pp. 194-214.

- Feldman, Richard. (2009). Evidentialism, higher-order evidence, and disagreement. *Episteme*, 6(3), 294-312.
- Fleisher, Will. (2021). How to Endorse Conciliationism. *Synthese*, 198(10), 9913-9939.
- Frances, Bryan, & Matheson, Jonathan. (2018). Disagreement. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2018 Edition). <https://plato.stanford.edu/archives/spr2018/entries/disagreement/>.
- Horowitz, Sophie. (2014). Epistemic Akrasia. *Noûs*, 48(4), 718-744.
- Kelly, Thomas. (2005). The Epistemic Significance of Disagreement. *Oxford Studies in Epistemology*, 1(167-196).
- Kelly, Thomas. (2010). Peer Disagreement and Higher Order Evidence. In R. Feldman and T. Warfield (Eds.), *Disagreement*. Oxford: Oxford University Press.
- King, Nathan. L. (2012). Disagreement: What's the Problem? Or a Good Peer is Hard to Find. *Philosophy and Phenomenological Research*, 85(2), 249-272.
- Knoks, Aleks. (2022). Conciliatory Views, Higher-order Disagreements, and Defeasible Logic. *Synthese*, 200(2), 173.
- Knoks, Aleks. (2021). Conciliatory Reasoning, Self-defeat, and Abstract Argumentation. *The Review of Symbolic Logic*, 1-48.
- Littlejohn, Clayton. (2020). Should we be Dogmatically Conciliatory?. *Philosophical Studies*, 177(5), 1381-1398.
- Littlejohn, Clayton. (2014). A Note Concerning Conciliationism and Self-defeat: A Reply to Matheson. *Social Epistemology Review and Reply Collective*, 3(12), 104–112.
- Matheson, Jonathan. (2009). Conciliatory Views of Disagreement and Higher-order Evidence. *Episteme*, 6(3), 269-279.
- Matheson, Jonathan. (2015). Epistemic Norms and Self-defeat: A Reply to Littlejohn. *Social Epistemology Review and Reply Collective*, 4(2), 26–32.
- Matheson, Jonathan. (2014). Are Conciliatory Views of Disagreement Self-Defeating?, *Social Epistemology: A Journal of Knowledge, Culture and Policy*, DOI: 10.1080/02691728.2014.907833
- Mulligan, Thomas. (2015). Disagreement Peerhood, and Three Paradoxes of Conciliationism. *Synthese*, 192(1), 67-78.

- Mulligan, Thomas. (2021). The Epistemology of Disagreement: Why not Bayesianism?. *Episteme*, 18(4), 587-602.
- Oppy, Graham. (2010). Disagreement. *International Journal for Philosophy of Religion*, 68(1-3), 183-199.
- Pittard, John. (2015). Resolute Conciliationism. *The Philosophical Quarterly*, 65(260), 442-463.
- Pittard, John. (2019). Fundamental Disagreements and the Limits of Instrumentalism. *Synthese*, 196(12), 5009-5038.
- Plantinga, Alvin. (2000). Pluralism: A Defense of Religious Exclusivism. In P. L. Quinn & K. Meeker (Eds.), *The Philosophical Challenge of Religious Diversity* (172-192.). New York: Oxford University Press.
- Smithies, Declan. (2019). *The Epistemic Role Of Consciousness*. Oxford University Press.
- Titelbaum, Michael. (2015) Rationality's Fixed Point (Or: In Defense of Right Reason), *Oxford Studies in Epistemology*, 5, 253–94.
- Wedgwood, Ralph (2010). The Moral Evil Demons. In Richard Feldman & Ted A. Warfield (Eds.), *Disagreement*. Oxford University Press.
- Whiting, Daniel. (2020). Higher-Order Evidence. *Analysis*, 80(4), 789-807.