

How to disagree about how to disagree*

Adam Elga

October 31, 2007[†]

Abstract

When one encounters disagreement about the truth of a factual claim from a trusted advisor who has access to all of one's evidence, should that move one in the direction of the advisor's view? Conciliatory views on disagreement say "yes, at least a little." Such views are extremely natural, but they can give incoherent advice when the issue under dispute is disagreement itself. So conciliatory views stand refuted. But despite first appearances, this makes no trouble for *partly* conciliatory views: views that recommend giving ground in the face of disagreement about many matters, but not about disagreement itself.

1 Introduction

Suppose that you and a friend independently evaluate a factual claim, based on the same relevant evidence and arguments. You become confident that the claim is true. But then you find out that your friend—whose judgment you respect—has become just as confident that the claim is false. Should that news at all reduce your confidence in the disputed claim?

Conciliatory views on disagreement answer "yes." According to such views, finding out that a respected advisor disagrees with one should move one at least a little in the direction of the advisor's view. And it should do so regardless of the subject matter under dispute. Conciliatory views are extremely natural and appealing. But they seem to run into trouble when the topic under dispute is disagreement itself. Can conciliatory views accommodate disagreement about disagreement? And if not, what does this show about what view on disagreement we should adopt instead?

*Thanks to Agustín Rayo, Delia Graff Fara, John Collins, Ted Sider, Brian Weatherston, David Enoch, the Corridor group, and an audience at the University of Wisconsin-Madison.

[†]To appear in Richard Feldman and Ted Warfield (eds.) *Disagreement*, forthcoming from Oxford University Press.

I will consider two arguments that conciliatory views cannot accommodate disagreement about disagreement. Though the first argument fails, the second argument succeeds. So conciliatory views are unacceptable. But the considerations that show this make no trouble for views that are partly conciliatory: views that recommend compromise in the face of disagreement about many matters, but not about disagreement itself.

2 *First argument against conciliatory views: repeated disagreements with the stubborn*

Can conciliatory views accommodate disagreement about disagreement? Here is a reason to think not.¹ Suppose that you and your friend disagree about the right response to disagreement. You have a conciliatory view, but you realize that your friend has the **stubborn view**, according to which disagreement is never cause for changing one's view on a disputed issue. It can sometimes seem as though your conciliatory nature dooms you to conceding everything to your stubborn friend, given enough discussion. Here is a representative scenario:

You think it will rain tomorrow, and your friend thinks it won't. (Here and henceforth I assume that you respect the opinions of all of your friends, and that you and your friends have the same evidence relevant to contested issues.) In response to the disagreement, you are conciliatory: you reduce your confidence that it will rain. But your friend is stubborn: he remains completely unmoved.

After this first stage, a (slightly less extreme) disagreement about the weather remains. Again you are conciliatory, and further reduce your confidence that it will rain. And again, your friend stands fast.

Disagreement still remains. You reduce your confidence a third time, and so on. As the discussion continues, you get pushed arbitrarily close to completely adopting your friend's view on whether it will rain.

¹I haven't seen this objection in print (though see Weatherson (2007, 8)), but have encountered it repeatedly in conversation. It deserves to be put to rest.

In this case, it looks as though your conciliatory nature commits you to conceding increasingly much, the more times you pool opinions with your stubborn friend. And this looks to be a general phenomenon. If so, that counts against conciliatory views on disagreement. For it is implausible that one should be required to give so much ground to an advisor just because the advisor is stubborn. A similar difficulty arises in the case of advisors who are not completely stubborn, but who have a policy of conceding very little in cases of disagreement.

That is the first argument against conciliatory views on disagreement.

3 *Reply: conciliatory folk needn't concede everything to stubborn folk*

Here is a reply: Sensible conciliatory views do not entail that one should concede arbitrarily much to stubborn advisors.

To see why not, imagine a cluster of advisors who you know to exhibit an extreme form of groupthink: they always end up agreeing with one another. Now, you may well respect the opinions of that group. So you may well be moved if you find out that one of them disagrees with you about a particular issue. But suppose that you then find out that another member of the group also disagrees with you about that issue. That news does not call for any additional change in your view. For you knew in advance that the group members all think alike. So hearing the second dissenting opinion gives you no real new information.

In contrast, suppose that you receive an additional dissenting opinion from an advisor who formed her opinions completely independently from your first advisor. In that case, the second dissenting opinion does call for additional caution. The difference is that in this case you didn't know in advance what conclusion the second advisor would reach.

The general point is that an additional outside opinion should move one only to the extent that one counts it as independent from opinions one has already taken into account.² The above example illustrates the most extreme version of this point: When one knows with certainty in advance what an advisor thinks, hearing that advisor's opinion should have no impact. But the point also holds in less extreme cases. For example, suppose that two of your friends *almost* always think alike. Then hearing that the first friend disagrees with you should have a big impact on your opinion.

²Compare Kelly (2007).

But suppose that you later learn that the second friend endorses the judgment of the first. That news should only have a tiny additional impact on your opinion.

The above independence point is completely uncontroversial, and every sensible view on disagreement should accommodate it.³ Furthermore, conciliatory views on disagreement face no special difficulties in doing so.

Now return to the case in which you disagree about the weather with a stubborn friend. When you find out about the initial disagreement, you should indeed be significantly moved in the direction of your friend's view. But at the second stage, news of the disagreement should not move you at all. The reason is the same as in the groupthink case: Since you knew in advance about your friend's stubborn nature, his continued disagreement provides you with no additional news. Putting things another way: You count his opinion at the first stage of the dispute as completely correlated with his opinion at subsequent stages. As a result, a sensible conciliatory view will counsel you to remain unmoved at the second and subsequent stages.

A similar analysis applies in the case of an advisor who is not completely stubborn, but who has a known policy of conceding very little in cases of disagreement. The initial disagreement of such an advisor should have a big impact on your opinion. But when the advisor keeps putting forward the same view in subsequent disagreements, that should have little or no additional impact.

Moral: Sensible conciliatory views do not require one to concede everything to stubborn advisors. That answers the argument.

³For example, according to the Equal Weight View (Elga 2007) it is a constraint on rationality that one's probability in a disputed claim match one's prior probability in the claim, conditional on what one has learned about the circumstances of the disagreement (see Elga (2007, footnote 26)). But when one is certain in advance what an advisor's reaction to the claim will be, that prior conditional probability will equal one's prior unconditional probability in the claim. So the Equal Weight View is consistent with the above observation about additional opinions (that hearing an additional opinion should only move one to the extent that one counts it as independent of information one has already taken into account).

4 *Second argument against conciliatory views: such views undermine themselves*

Next argument.

Just as people disagree about politics and the weather, so too people disagree about the right response to disagreement. For example, people disagree about whether a conciliatory view on disagreement is right. So a view on disagreement should offer advice on how to respond to disagreement about disagreement. But conciliatory views on disagreement run into trouble in offering such advice.

The trouble is this: In many situations involving disagreement about disagreement, conciliatory views call for their own rejection. But it is incoherent for a view on disagreement to call for its own rejection. So conciliatory views on disagreement are incoherent. That is the argument.⁴

To see why conciliatory views sometimes call for their own rejection, consider an example. Suppose that you have a conciliatory view on disagreement, but you find out that your respected friend disagrees. He has arrived at a competing view (about disagreement), and tells you all about it. If your conciliatory view is correct, you should change your view. You should be pulled part way toward thinking that your friend is right. In other words, your view on disagreement requires you to give up your view on disagreement.

One might try to avoid this result by adding a special restriction to one's conciliatory view. For example, might say that one should *in general* be moved by disagreement, but not when the disputed topic is disagreement itself. But such a restriction seems objectionably arbitrary and ad hoc.⁵ If one should be sensitive to disagreement about so many other matters, then why not about disagreement, too? (Certainly not because disagreement is an easy or uncontroversial topic, as the existence of this volume attests.)

So: Conciliatory views on disagreement sometimes call for their own rejection. The next section explains why views on disagreement that call for their own rejection are incoherent. It will follow that conciliatory views

⁴I first learned of this objection from an unpublished early draft of Kelly (2005), which discusses the objection without endorsing it. Weatherson (2007) has independently raised and developed an objection of this kind.

⁵Disclosure: it will later emerge that a similar restriction *should* be imposed. But it will take real work to explain away the seeming arbitrariness of doing so.

on disagreement are incoherent.

5 *Self-undermining views are incoherent*

Why is it incoherent for a view on disagreement to call for its own rejection? To see why, notice that one's view on disagreement is part of one's *inductive method*: one's fundamental method for taking evidence into account. An inductive method offers recommendations on what to believe based on one's course of experience. Given a course of experience, an inductive method says what one should believe about various topics: the weather, who will win the next election, and so on. It even says how a given course of experience bears on the question: Which inductive method should one use?

Now suppose that one's view on disagreement sometimes calls for its own rejection. Then one's inductive method also sometimes calls for its own rejection. For one's view on disagreement is part of one's inductive method.⁶ So in order to show that self-undermining views on disagreement are incoherent, it is enough to show that self-undermining inductive methods are incoherent.

That is best illustrated by the following example.⁷

The magazine *Consumer Reports* rates appliances, and gives recommendations on which ones to buy. But pretend that in addition to rating appliances, *Consumer Reports* also rates and recommends consumer ratings magazines. Then it cannot coherently recommend a competing magazine over itself. (By a "competing magazine" I mean a magazine that offers contrary appliance recommendations.)

To see why not, consider an example. Suppose that *Consumer Reports* says, "Buy only Toaster X," while *Smart Shopper* says, "Buy only Toaster Y." And suppose that *Consumer Reports* also says, "*Consumer Reports* is worthless. *Smart Shopper* magazine is the ratings magazine to follow." Then *Consumer Reports* offers inconsistent advice about toasters. For on

⁶More slowly: Suppose that one has view V on disagreement, and suppose that one has inductive method M . Then view V must be part of method M . So if (given a particular course of experience) view V says to reject view V , M must (given that same course of experience) say to reject view V . That is because M says everything V says. But to reject view V is to reject M , since V is part of M . So M says to reject M . So if V is self-undermining, then M is also self-undermining.

⁷The *Consumer Reports* analogy is adapted from Lewis (1971, 55).

the one hand, it says directly to buy only Toaster X. But on the other hand, it also says to trust *Smart Shopper*, which says to buy only Toaster Y. And it is impossible to follow both pieces of advice.

In other words:

1. *Consumer Reports* says: "Buy only toaster X."
2. *Smart Shopper* says: "Buy only toaster Y."
3. *Consumer Reports* says: "Follow the advice of Smart Shopper."

Given what *Smart Shopper* says about toasters, items 1 and 3 offer conflicting advice. So *Consumer Reports* gives conflicting advice about toasters. And a similar conflict arises in any case in which *Consumer Reports* recommends a competing magazine over itself.

Moral: No consumer rating magazine can coherently recommend a competing magazine over itself. For the same reason, no inductive method can coherently recommend a competing inductive method over itself. Let me explain, using an argument adapted from Field (2000, 131).

Just as a consumer ratings magazine tells one how to shop, an inductive method tells one how to respond to various courses of experience. An inductive method says something of the form, "Given course of experience E_1 , adopt *such-and-such* belief state. Given course of experience E_2 , adopt *so-and-so* belief state. Given course of experience E_3 , adopt *blah-blah-blah* belief state, ..." In other words, an inductive method puts forward a rule for responding to possible courses of experience.

One small bit of terminology: Given an initial course of experience, let us say that two inductive methods are *competitors* (and that each is a *competing* method to the other) if they offer contrary recommendations about how to respond to some possible subsequent experience.

Now: It is incoherent for an inductive method to recommend two incompatible responses to a single course of experience. But that is exactly what a method does if it ever recommends a competing method over itself.

For example, suppose that inductive methods M and N offer contrary advice on how to respond to the course of experience "See lightning, then see a rainbow." In particular, suppose:

1. Method M says: "In response to seeing lightning and then a rainbow, adopt belief state X."

2. Method *N* says: “In response to seeing lightning and then a rainbow, adopt belief state *Y*.”

(Assume that it is impossible to adopt both belief states *X* and *Y*.) But also suppose that *M* sometimes calls for its own rejection:

1. Method *M* says: “In response to seeing lightning, stop following method *M* and start following method *N*.”

Then method *M* offers inconsistent advice. On the one hand, it directly recommends belief state *X* in response to seeing lightning and then a rainbow. But on the other hand, it also says that seeing lightning should make one follow method *N*, which recommends belief state *Y* in response to seeing lightning and then a rainbow. And it is impossible to follow both pieces of advice. So method *M* gives incoherent advice about how to respond to seeing lightning then a rainbow. And a similar conflict arises in any case in which an inductive method recommends a competing method over itself.⁸

So: Just as a consumer ratings magazine cannot consistently recommend a competing magazine, an inductive method cannot consistently recommend a competing method. In other words, self-undermining inductive methods are incoherent. It follows that conciliatory views on disagreement are incoherent. Call this the *self-undermining problem*.

Bottom line: The self-undermining problem shows that conciliatory views on disagreement should be rejected.

⁸It might be thought that some conciliatory views on disagreement avoid this problem because they do not entirely call for their own rejection. Rather, they merely call for their own *partial* rejection. For example, consider a case in which someone with a conciliatory view—call it *C*—learns about a respected friend’s competing view of disagreement—call it *D*. The conciliatory view need not say in this case, “Reject *C* and adopt *D*.” Instead it might say, “Become uncertain as to whether *C* or *D* is the right view on disagreement.”

But even views on disagreement that call for their own partial rejection are incoherent. For notice that when one shifts one’s view about the right way to respond to disagreement, one should correspondingly shift the way one responds to subsequent disagreements. In particular, when the above subject shifts his confidence away from view *C* and towards view *D*, that should correspondingly change the inductive method he implements. It won’t be as dramatic a change as if he had become completely converted to view *D*, but it will be a change nonetheless. In other words, even in this sort of case, view *C* calls for a change in inductive method. And for certain choices of view *D*, view *C* calls for a change to a *competing* inductive method. But now the argument in the main text applies. For that argument applies to any inductive method that recommends a competing method over itself.

6 *Reply to the self-undermining problem for conciliatory views*

There is no good reply. Conciliatory views stand refuted.

7 *If conciliatory views are wrong, should we adopt an uncompromising view instead?*

Conciliatory views get into trouble because they require one to be conciliatory about *absolutely everything*, even their own correctness. But we have seen that it is incoherent to be conciliatory about absolutely everything.⁹ So conciliatory views are no good. What view should we adopt instead? We might adopt a view that is conciliatory about many matters, but not about disagreement itself. But as noted before, such views seem to require arbitrary and ad hoc restrictions.

Alternatively, we might adopt a view that avoids the self-undermining problem without imposing special restrictions. We have already seen one such view: the stubborn view. The stubborn view avoids the self-undermining problem because according to the stubborn view, disagreement about disagreement shouldn't at all affect one's views on disagreement. So there is no threat of the stubborn view ever calling for its own rejection.

A more plausible view that also avoids trouble in cases of disagreement about disagreement is the **right-reasons view**.¹⁰ The right-reasons view is best explained with an example: Dee and Dum independently assess a claim, based on the same batch of evidence E. When they later find out that they came to opposite conclusions, how should they react? According to

⁹So it is a good thing that some authors who defend conciliation in a great range of cases stop short of advocating it across the board. For example, Feldman (2007) gives arguments that favor suspending judgment in symmetric cases of disagreement. But he only claims that suspension of judgment is required "At least for some range of hard cases." (Feldman 2007, 212). Similarly Christensen (2007, 189) limits his endorsement of conciliation to a restricted range of cases: "I shall argue that in a *great many cases* [of peer disagreement] of the sort van Inwagen and others seem to have in mind, I should change my degree of confidence significantly toward that of my friend [emphasis added]." Even the Equal Weight View (Elga 2007) falls short of requiring conciliation about all topics. For that view takes the form of a constraint on conditional probabilities (see footnote 3). As a result, the view is compatible with thinking that agents should have probability 1 in certain propositions, and that no news of disagreement should reduce that probability.

¹⁰The right-reasons view is a simplified version of the view defended in (Kelly 2005, 180).

the right-reasons view, that depends on what conclusion evidence E in fact supports. For example, suppose that E supports Dee's conclusion. Then in reaction to the disagreement, Dee should stick to that conclusion, and Dum should switch to it.

More generally, the right-reasons view says that in the face of disagreement, one should adopt whatever view one's original evidence in fact supports. Here one's "original evidence" is the evidence that one had before finding out about anyone else's conclusions.

The right-reasons view has no special trouble accommodating disagreement about disagreement. For example: Suppose that your evidence strongly supports the right-reasons view, and that as a result you hold the right-reasons view. And suppose that you learn that a respected advisor holds a different view about disagreement. According to the right-reasons view, this should not at all weaken your confidence in the right-reasons view. In other words, according to the right-reasons view, encountering disagreement about disagreement in this case should have no effect at all on your opinions about disagreement. Other cases are similar.

Both the stubborn view and the right-reasons views are uncompromising in the following sense: Each entails that if one has correctly judged how one's original evidence bears on a claim, then just finding out that a respected advisor disagrees should not at all change one's confidence in the claim. In other words, while conciliatory views say that disagreement should *always* move one, these uncompromising views say that disagreement should *never* do so (provided that one has correctly responded to one's original evidence).

So: We have seen two ways that a view on disagreement can coherently handle cases of disagreement about disagreement. The view can be *partially conciliatory*, and say that one should be moved by disagreement about some subject matters, but not about disagreement itself. Or it can be *uncompromising* and say that one should not be moved by disagreement about any topic (provided that one has correctly responded to one's original evidence). But partially conciliatory views seem to require arbitrary and ad hoc restrictions. So the undermining problem seems to favor adopting an uncompromising view (Weatherson 2007, 8).

But this is an illusion. It is not at all arbitrary for a view on disagreement to treat disagreement about disagreement in a special way. So the self-undermining problem is no evidence for uncompromising views about disagreement. Here is why.

8 *The source of the self-undermining problem*

It looks arbitrary for a view to recommend that one be conciliatory about most matters, but not about disagreement itself. But in fact no arbitrariness is required. For the discussion of *Consumer Reports* and inductive methods shows that it is in the nature of giving consistent advice that one's advice be dogmatic with respect to its own correctness. And views on disagreement give advice on how to respond to evidence. So in order to be consistent, views on disagreement must be dogmatic with respect to their own correctness.

In other words, the real reason for constraining conciliatory views is not specific to disagreement. Rather, the real reason is a completely general constraint that applies to any fundamental¹¹ policy, rule, or method. In order to be consistent, a fundamental policy, rule or method must be dogmatic with respect to its own correctness. This general constraint provides independent motivation for a view on disagreement to treat disagreement about disagreement in a special way. So partly conciliatory views need no ad hoc restrictions in order to avoid the self-undermining problem. They only need restrictions that are independently motivated.

Let me illustrate the point with a *Consumer Reports* example. Suppose that for twenty-eight years in a row, *Consumer Reports* rates itself as the #1 consumer ratings magazine. A picky reader might complain to the editors:

You are evenhanded and rigorous when rating toasters and cars. But you obviously have an ad hoc exception to your standards for consumer magazines. You always rate yourself #1! Please apply your rigorous standards across the board in the future.

This complaint has no force. The editors should reply:

To put forward our recommendations about toasters and cars *is* to put them forward as good recommendations. And we can't consistently do that while also claiming that contrary recommendations are superior. So our always rating ourselves #1 does not result from an arbitrary or ad hoc exception to our

¹¹A fundamental method is one whose application is not governed or evaluated by any other method. See Field (2000, Appendix).

standards. We are forced to rate ourselves #1 in order to be consistent with our other ratings.

The same point holds for views of disagreement. Just as *Consumer Reports* has good independent motivation to avoid recommending a competing magazine, so too a view on disagreement has good independent motivation to avoid calling for its own rejection. In particular, partly conciliatory views have good independent motivation for treating the case of disagreement about disagreement differently from cases of, say, disagreement about the weather.

Bottom line: Partly conciliatory views need no ad hoc restrictions to avoid the self-undermining problem. So the self-undermining problem does not favor uncompromising views over partly conciliatory ones. So even though considerations arising from disagreement about disagreement refute views that are conciliatory about *every* topic, they are no evidence against views that are conciliatory about a great many topics.

References

David Christensen. Epistemology of disagreement: the good news. *Philosophical Review*, 116:187–217, 2007.

Adam Elga. Reflection and disagreement. *Noûs*, 41:478–502, 2007.

Richard Feldman. Reasonable religious disagreements. In Louise M. Antony, editor, *Philosophers without Gods: Meditations on Atheism and the Secular Life*. Oxford University Press, 2007.

Hartry Field. Apriority as an evaluative notion. In Paul Boghossian and Christopher Peacocke, editors, *New Essays on the A Priori*, chapter 6, pages 117–149. Oxford University Press, 2000.

Thomas Kelly. The epistemic significance of disagreement. In John Hawthorne and Tamar Gendler, editors, *Oxford studies in epistemology*, volume 1. Oxford University Press, Oxford, 2005.

Thomas Kelly. Peer disagreement and higher order evidence. In Rich Feldman and Fritz Warfield, editors, *Disagreement*. Oxford University Press, 2007. Forthcoming.

David Lewis. Immodest inductive methods. *Philosophy of Science*, 38(1): 54–63, 1971.

Brian Weatherson. Disagreeing about disagreement. Manuscript, January 2007. URL <http://brian.weatherson.org/DaD.pdf>.