

[Final version in The Epistemology of Disagreement: New Essays, edited by David Christensen and Jennifer Lackey, OUP 2013. Please feel free to email me for a pdf of the published version: david.christensen@brown.edu.]

## **Epistemic Modesty Defended<sup>1</sup>**

David Christensen

Many recent writers have embraced one version or another of the thought that the disagreement of equally informed, equally skillful thinkers can (in at least some circumstances), require a rational agent to revise her beliefs (to at least some extent)—even if her original assessment of the common evidence was correct. There has, of course, been much disagreement as to the amount of revision required in certain cases, and as to the theoretical underpinnings of the required revisions. But a common thread uniting all these views is the recognition that we may make mistakes in assessing evidence; that the disagreement of others who have assessed the same evidence differently provides at least some reason to suspect that we have in fact made such a mistake; and that reason to suspect that we've made a mistake in assessing the evidence is often also reason to be less confident in the conclusion we initially came to. The rationale for revision, then, expresses a certain kind of epistemic modesty.

So far, this may seem like little more than epistemic common sense. But it turns out that the sort of modesty in question has some puzzling consequences. The consequences have come out mainly in discussions of positions advocating highly conciliatory responses to disagreement. On such positions, if I hold some philosophical view, for example, and find myself in disagreement with other philosophers—philosophers familiar with all the same arguments, and

---

<sup>1</sup> Versions of this paper were presented at the Collège de France, University of Oxford and The Ohio State University; thanks to all three audiences for helpful discussion. I'd also like to thank Stew Cohen, Adam Elga, Sophie Horowitz, Jennifer Lackey, John Pittard, Josh Schechter, Jonathan Vogel and the participants in my graduate seminar for valuable discussion and/or comments on earlier drafts.

possessing philosophical skills equivalent to mine—I often should become much less confident in my philosophical view, perhaps, in categorical-belief terms, withholding belief on the topic.<sup>2</sup>

Suppose I hold such a view (call it CV, for Conciliatory View<sup>3</sup>), and that I practice what I preach. So, for example, when I think about the arguments directly relevant to a certain version of mentalist Internalism about epistemic justification, it seems very likely to me that it's true. But in response to the disagreement of epistemologists I respect, I become much less confident in Internalism. Now as it turns out, I'm also aware of the current controversy about disagreement, and know that a number of epistemologists reject CV in favor of positions toward the “steadfast” end of the spectrum: they hold that one may (often, at least in large measure) maintain one's confidence in one's initial beliefs despite knowledge of disagreement by those who seem, independent of the disagreement, to be as well positioned as oneself to arrive at accurate views on the disputed matter. I also quite reasonably respect epistemologists who hold steadfast views and reject CV. Insofar as I practice what I preach, it seems that CV requires me to become much less confident in CV as well.

This puts the advocate of CV in a situation that's puzzling in a number of ways. For one thing, it would seem that, in the present epistemological climate, at least, CV has the property that one cannot rationally believe it (at least very strongly), even if it's true. But this in itself isn't obviously mysterious or deeply problematic. After all, there would seem to be other situations—ones in which all epistemologists accept CV, for instance—in which one could rationally believe in CV. So CV isn't obviously intrinsically impossible to believe rationally. The present situation might, for all we've seen so far, simply be the sort of situation we confront on all kinds of topics all the time: one in which the truth on some matter is not rationally believable, because our evidence, misleadingly, points away from it.

---

<sup>2</sup> Of course, the description of the case needs filling out in various ways. I'm assuming that there is not a huge imbalance in the number of philosophers holding the different views, that there is not a preponderance of more highly skilled philosophers on one side, that I have reason to think that the stated views of the philosophers involved reflect their ordinary and honest appraisal of the arguments, not bizarre brainwashing, joking, etc.

<sup>3</sup> An often-used term to describe certain views of this sort is “Equal Weight View”. This term was used by Elga (2007) to describe his own view, but has been used by various writers, not all of whom use it the same way. My term is intended to be more general, and not to suggest limitation to, e.g., Elga's version of the position. The term “conciliatory,” is actually also taken from Elga (see his 2010).

But there are a couple of more serious worries in this general neighborhood. In the next section, I'll look at the worry that the self-undermining character of CV makes it impossible to maintain any stable view of disagreement that doesn't reject CV completely. In the sections that follow, I'll look at an argument which uses self-undermining to show that CV, and many other principles expressing epistemic modesty, are inconsistent, and thus must be rejected. I will argue that, despite these difficulties, epistemic modesty can be defended.

### **1. Self-undermining and Instability**

The first worry emerges when we think in more detail about how someone who is initially convinced by conciliationist arguments should react to the anti-CV beliefs of philosophers she respects. We can see the problem even in a simple case where we abstract from the wider debate and consider just Connie, a conciliationist, and Steadman, who holds a steadfast view. Suppose that Connie and Steadman are unaware of the wider debate. Connie, in her study, thinks hard about the issue and becomes highly confident of CV—say, she reaches credence .99—and goes to tell Steadman. He tells her that he, too, has been thinking hard about disagreement, but that he has become equally confident of a steadfast view SV. They discuss the issue thoroughly, each offering his or her own arguments, and conscientiously attending to the arguments of the other. At the end of the day, unfortunately, they still disagree just as strongly. Let us assume, for the sake of argument, that Connie's original take on these arguments is correct; these arguments do in fact support CV, and in fact Steadman has misjudged them.

Connie, back in her study, reflects on the day. She considers Steadman her philosophical equal; in fact, she has strong reason to believe that Steadman is as likely to get things right in philosophy as she is. And she has every reason to think that Steadman's general parity in philosophical reasoning would extend to thinking about the epistemology of disagreement (at least, insofar as she puts aside the fact that he's arrived at SV, which seems wrong to her). So, good conciliationist that she has become, Connie reduces her confidence in CV dramatically. Say, for the sake of simplicity, that she now has about equal confidence in CV and SV.

But there seems to be something unsatisfactory about where Connie has ended up. From her perspective when she emerged from her study, she did the right thing in fully conciliating with Steadman. But from her present perspective, there's only about half a chance that CV is

correct (in which case she reacted correctly to the fact of Steadman's disagreement). There's also about half a chance that SV is correct—in which case she should ignore disagreement and maintain the belief that's supported by the original arguments; that view, by her lights (and in fact) is CV. What should her reaction be to this new uncertainty about the correct rules for belief-management?

A natural suggestion is something like this: Insofar as she divides her credence between two different rules for belief-management, and those rules dictate different credences for some proposition, she should adopt a credence in between what the two rules recommend: a mixture of the recommended credences. If she's much more confident of the first of the rules, her new credence in the relevant proposition should be closer to that recommended by the first rule. In cases such as the present one, where she has about equal confidence in the two rules, she should adopt a credence in the disputed proposition that's about halfway between the credences recommended by the two rules.

Let us suppose for the sake of argument that something along the lines of this natural thought is correct, and that Connie sees this. It seems that she should decide that her present credence in CV is too low—that she went too far in her initial compromise with Steadman. For while CV recommends her present credence of about .5, SV recommends that she have .99 credence. If she mixes those recommendations equally, she'll arrive at around .75 credence in CV.

But suppose she does this. It seems that she's still in trouble. For now she again thinks that it is much more likely that she should be conciliatory than that she should be steadfast. So it should now seem to her that she hasn't conciliated enough. Applying the common-sense thought again, she should readjust her credence in SV to a mixture of about 75% the conciliatory compromise between her original assessment and Steadman's, and 25% the SV-mandated credence produced by her original assessment of the arguments. This will land her in the vicinity of having .62 credence in CV. And the process continues...<sup>4</sup>

---

<sup>4</sup> The worry that a defender of CV, confronted with disagreement about CV, will not be able to form a stable credence in CV is due to Weatherson (2007), reprinted as part of his contribution to this volume. Weatherson uses a somewhat more complex argument for his conclusion, but I think that the argument sketched in the text, which depends on the same sort of assumption about how beliefs about the correct rule for belief-management interact with beliefs in general, makes the same point.

A couple of different questions arise at this point. One of them is whether there is any end to the series of adjustments that Connie has begun. A blog post by Matt Weiner<sup>5</sup> suggests that there is. We might formalize the common-sense thought somewhat as follows, at least for the special case where an agent's credence about the right epistemic rules is divided between rule A and rule B (let  $A_p$  be the credence in  $p$  that rule A would recommend, and  $B_p$  be the credence in  $p$  that rule B would recommend, and let  $A$  and  $B$  stand, respectively, for the claims that rules A and B are correct):

$$(*) \text{ cr}(p) = \text{cr}(A) \cdot A_p + \text{cr}(B) \cdot B_p$$

In other words, the agent's credence in  $p$  should be a weighted average of the credences rule A and rule B would recommend for her, where the weighting is determined by how much credence she has in the correctness of rules A and B.<sup>6</sup>

Weiner assumes that the agent in question begins by having credence 1 in CV, and her acknowledged peer begins with credence 1 in SV.<sup>7</sup> Is there a stable position for the agent to take, once she applies (\*) to her view about disagreement? Given the description of the case, let us suppose that CV recommends that she split the difference, arriving at .5 credence in CV, and that SV recommends that she retain her original view: full credence in CV. Putting these recommendations into (\*) above, we get:

$$\text{cr}(CV) = \text{cr}(CV) \cdot .5 + (1 - \text{cr}(CV)) \cdot 1.$$

Weiner points out that the equation balances when  $\text{cr}(CV)$  is  $2/3$ . In other words, if the agent adopts  $2/3$  credence in CV, her view is consistent with (\*), which represents our common-sense

---

<sup>5</sup> Weiner (2007); this is a response to Weatherson (2007). Weiner's post centers around an example more like the one discussed here.

<sup>6</sup> Very similar principles are put forth by Weatherson and Weiner in their discussions of instability.

<sup>7</sup> There may be problems with this if credence 1 is interpreted in standard ways. But for present purposes, let us put this aside; the example works equally well with credences near, but not at, the extremes.

thought about how to react to uncertainty about epistemic rules. So perhaps this is where Connie should end up in our beginning example. Weiner also indicates how this solution can generalize to other cases of the same sort, but where the agents in question begin with different degrees of credence in CV and SV. Thus defenders of CV may hope that the view does not, after all, lead to problematic instability, at least in Connie's sort of case.

However, it is worth noticing that there is something quite odd about the above discussion. Consider how Connie should think about her credence in CV, supposing that she's obeying (\*), and has settled on  $2/3$  credence in CV and  $1/3$  credence in SV. (\*) captures the thought that Connie should weight the recommendations of the two rules according to how likely she thinks they are to be correct—that is, to correctly describe rational belief. But while this seems plausible when put in the abstract, its plausibility is considerably strained when one considers how Connie should regard her credences *in the correctness of those epistemic rules*.

(\*) would have Connie balance two credence-components: the credence recommended by CV, and the credence recommended by SV. As we've described the story so far, CV recommends a much lower credence in CV than SV recommends. Thus to the extent that Connie favors the CV-recommended credence, she'll be less confident in CV. But she's supposed to weight these factors by her credences in the respective correctness of two rules—by how likely she takes them to be. Thus to the extent that she thinks CV is correct, she'll be led to lower her credence in CV; and to the extent that she thinks SV is correct (and thus CV is incorrect), she'll be led to raise her credence in CV! And this fact is perfectly transparent to Connie. It's not at all clear that, in this situation, Connie's following (\*) would be reasonable.

To put the point a slightly different way, consider how Connie should think about her own credence. It would be natural for her to think through her own complying with (\*) somewhat as follows:

“Well, suppose that CV is true. In that case, I shouldn't be very confident of it. And CV is probably correct. So I shouldn't be too confident in it.”

“Suppose SV is correct. In that case, I should stick to my guns and be highly confident that SV is false. There's a decent chance that SV is correct. So I should have a decent amount of credence that it's incorrect.”

Shouldn't Connie arrange her beliefs instead so that, to the extent that CV is likely to be true, she has *high* credence in CV? There seems to be something fundamentally incoherent in Connie's reasoning in a way that manifestly reverses this relationship.

It's worth noting that this is not a difficulty with Weiner's *solution* to the instability problem. The difficulty arises from the common-sense thought which generated the instability in the first place, as formulated by (\*). CV itself just says that Connie should reduce her confidence in a certain proposition (that CV is correct) after talking to Steadman. This in itself is not sufficient to generate any instability at all. But when Connie's reduced confidence in this proposition is taken to require her to change her epistemic procedures in the way recommended by (\*), then the difficulties arise. And while it seems extremely plausible that one's beliefs about things in general ought in some way to be sensitive to one's beliefs about the rules of rationality,<sup>8</sup> we've just seen that one seemingly natural way of implementing this plausible idea generates puzzling results.<sup>9</sup> And this puzzlement is independent of whether there is a stable way for Connie to comply with (\*).

So for the present, I'd like to leave off worrying about instability per se. But I want to keep one lesson from the above discussion in mind: The interaction between CV and (\*) suggests that the oddity of self-undermining epistemic principles is closely tied to questions about how credences about what credences are rational constrain credences in general. This connection will, I think, emerge more clearly in what follows.

## **2. Self-undermining and Inconsistency**

The second worry about CV is not centered around stability, but around whether it is intrinsically incoherent. The defender of CV, as noted above, might start out being quite sanguine about the

---

<sup>8</sup> Isn't that one of the main reasons people have given for being interested in epistemology?

<sup>9</sup> For discussion of a principle ("Rational Reflection") that is closely related to (\*), see Christensen (2011b), where more dramatic problems with the principle are developed. I should note that unpublished work by Adam Elga argues that these problems can be avoided by a revised principle, which still captures much of the intuitive appeal of Rational Reflection. Interestingly, I believe that Elga's principle may also avoid the instability problem described above. But it would not be appropriate for me to enter into the details here.

possibility that CV is not rationally believable at present, since plenty of true things are not rationally believable, given our current evidence. He might note that the opinions of others are important evidence, but that it's quite possible that the evidence in this case is misleading. Of course, he'll have to admit that, though he defends CV, he's not rationally confident in its truth. But why not go on defending it anyway, while others defend other views? After all, research goes best when different investigators pursue different lines simultaneously. Thus David Christensen (2009, 763) optimistically suggests: "Of course, we may still work hard at producing and disseminating arguments for the view, hoping to hasten thereby the day when epistemic conditions will brighten, consensus will blossom, and all will rationally and whole-heartedly embrace Conciliationism." But it seems to me that this degree of sanguinity underestimates the difficulty of the problem considerably.<sup>10</sup>

The reason is as follows: The sanguine line admits that CV, in the present epistemic circumstances, requires me to have only moderate credence in its own correctness. Plausibly, this means that I shouldn't completely follow CV in my present circumstances. But if that's right, then CV does not accurately describe how I should govern my beliefs in my present circumstances. But if CV were a correct general principle, one might argue, it would give correct direction in the present case. So it's not a correct principle.<sup>11</sup> Indeed, the argument suggests that even if conciliationists were tremendously successful in swaying philosophical opinion, so that the experts were uniformly confident that CV was correct, that *this* would be a case where the evidence provided by expert opinion turned out to be misleading!

A version of this problem is explored in detail by Adam Elga, who takes it to be an instance of a more general problem that occurs whenever an inductive method calls for its own rejection:<sup>12</sup>

---

<sup>10</sup> Thanks to Josh Schechter for helping me to realize this.

<sup>11</sup> Weatherson (2007) presses an argument very close to this.

<sup>12</sup> Elga attributes the general argument to Field (2000); a related argument occurs in Lewis (1971).



[I]t 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. (2010, 181-182)

The argument would seem to apply to CV as follows: Suppose that the direct philosophical arguments and evidence strongly support Internalism about epistemic justification, and that I assess the arguments and evidence correctly. CV, as noted above, will require that I still not be highly confident in Internalism, given that so many epistemologists I respect take the

arguments and evidence to support Externalism. Suppose that, on CV, I should move from .9 credence in Internalism (which is where I'd be on the basis of my considering the direct arguments alone), to .52 on the basis of disagreement. And suppose that I, having been convinced of CV, do move to .52 on Internalism. Then I'm confronted by disagreement about CV itself, and, again following CV, become much less confident of it. But now that I have serious doubts about whether CV is correct, I am more confident that I should pay less attention to disagreement. In light of that, it seems that I should now not compromise so much with others on Internalism, and should adopt a credence higher than .52. But by hypothesis, any credence other than .52 would violate CV. So following CV requires me to violate CV. So CV is inconsistent. Call the argument instantiated here the Inconsistency Argument against CV.

Elga takes this argument to apply to CV, or any complete inductive method that includes CV, since, as we've seen, a follower of CV is sometimes bound by CV to lose confidence in CV. It's worth noting that, on this way of seeing things, it is not the fact that CV recommends against itself *in the present circumstances* that causes the problem. Even if everyone happened to believe in CV, CV would still be incoherent, according to this argument. For it would still offer inconsistent recommendations for the possible situation in which many experts reject CV. Elga takes the problem to be decisive: "There is no good reply. Conciliatory views stand refuted." (2010, 182)

I take the sort of argument Elga presents to provide a very powerful challenge to CV. But before giving a final assessment of its force, I'd like to examine the argument more closely.

### **3. The Inconsistency Argument and Level-Connections**

One thing that seems to lie a bit below the surface of the argument as presented above is that it depends on some sort of level-connecting principle: some principle of rationality which connects credences about what credences are rational with credences in general. We saw in section 1 that the instability argument relied on a particular version of this idea. But the general idea seems implicit in the Inconsistency Argument as well.

To see this, note that CV is a rule for believing. It constrains my credences on controversial topics under certain conditions. The Inconsistency Argument begins by noting that CV will, in certain circumstances, apply to my credence in CV itself: it will prevent me from

being highly confident in the correctness of CV. But the argument then takes this to have implications for what I should believe about other things, such as Internalism. These implications do not, strictly speaking, follow from CV. If I lose confidence in CV (due to disagreement), this does not entail that I change my credence in Internalism. In fact, it would seem that I could obey CV completely, simply by maintaining my CV-mandated credence in Internalism, even after drastically lowering my credence in CV—the very principle that mandated that credence. So it seems that CV, taken neat, does not after all offer inconsistent advice.

To be sure, there would be something quite unattractive about making the doxastic move just envisioned. I would be regulating my credences in accordance with a principle whose correctness I seriously doubted—after all, CV provided my only reason for not being highly confident in Internalism. Can I rationally just ignore the fact that my lowered credence in Internalism was justified by a rule in which I no longer have confidence? There does seem to be something epistemically wrong with divorcing what it's rational for me to believe in general from what it's rational for me to believe about what's rational for me to believe. So my point here is not at all that dependence on such "level-connecting" intuitions deprives the Inconsistency Argument of force. The point is just that the considerable force the argument does have derives in part from a commitment to level-connecting.

Two other points seem worth making about this aspect of the argument. First, denying level-connections is not only intrinsically implausible; it also should be a particularly uncomfortable option for the defender of CV. As Weatherson points out, the intuitive plausibility of CV derives in large part from the appeal of some sort of level-connection. My reason for losing confidence in Internalism when I find out that others disagree derives from worrying that my response to the direct evidence and arguments was not, after all, the most rational one. I worry that I've incorrectly assessed the direct evidence and arguments. If this sort of worry about the rationality of my belief created no rational pressure for me to revise my beliefs, CV would be unmotivated.

The second point is that while the Inconsistency Argument depends on some sort of level-connection principle, it does not seem to depend on any very particular version of the idea, as does the Instability Argument. It seems to require that rational doubts about the correctness of a certain epistemic principle should weaken the extent to which that principle governs one's

beliefs in general. But it does not seem to require, e.g., that one believe in accordance with (\*). I take this to be a way in which the Inconsistency Argument is more powerful than the Instability Argument.<sup>13</sup>

#### 4. The Scope and Power of the Inconsistency Argument

Let us now look at a reason why one might be quite suspicious of the Inconsistency Argument. The argument is naturally raised against heavily conciliatory views of disagreement. But it seems clear that the argument, if sound, would have much wider application.

Consider first the spectrum of views on the rational reaction to disagreement. Many critics of strongly conciliatory views of disagreement have advocated positions on which one need not compromise much with apparent epistemic peers in many cases—say, because the reasons behind one’s original judgment on the disputed issue can do double duty and support the judgment that one’s apparent peer has made a mistake in evaluating the evidence relevant to that particular issue. Still, such moderately steadfast views typically concede that the disagreement of others has *some* tendency to affect one’s degree rational of confidence—just not nearly as strong a tendency as is claimed by more conciliatory views.<sup>14</sup> But it seems clear that such views are every bit as vulnerable to the Inconsistency Argument as are strongly conciliatory views. For insofar as disagreement has any power to reduce one’s rational credences in general, it will presumably have the power to reduce one’s rational credence in one’s moderately steadfast view. And insofar as that will require one to adopt credences different from those recommended by the moderately steadfast view, the Inconsistency Argument will apply.

In fact, it would seem that any view about disagreement short of absolute steadfastness is subject to the same problem. Consider:

---

<sup>13</sup> For example, I believe that the particular level-connecting principle Elga proposes as an improvement on Rational Reflection will enable the Inconsistency Argument, even if it does not enable the Instability Argument.

<sup>14</sup> Some examples of such views can be found in Kelly (2010), Lackey (2010, 2010a), Sosa (2010), and Wedgwood (2010). Even Kelly (2005), which advocates a strongly steadfast response to peer disagreement, acknowledges that disagreement of epistemic superiors—say those who one has reason to believe are less likely than oneself to be biased—calls for epistemic deference.

**Minimal Humility:** If I've thought about some complex issue P for ten minutes, and have decided P is true, and then find out that many people, most of them smarter than I, have thought long and hard about P, and have independently but unanimously decided that P is false, I should become less confident in P.

Clearly, Minimal Humility is subject to the Inconsistency Argument.<sup>15</sup> Indeed, as Bryan Frances points out,

If you accept *any* thesis T of the form 'If such-and-such conditions obtain with respect to one's belief in P, then one should withhold belief in P', then provided the antecedent of T can be true when  $P = T$ , then you have to deal with a version of the Self-Application Objection. The probability that some such principle T should be true gives us good grounds to think that there must be something amiss with the Self-Application Objection.<sup>16</sup>

Frances's abstract way of putting the point brings out that it is not just about the belief-undermining power of *disagreement*.<sup>17</sup> Suppose, for example, one has recently adopted a philosophical view on some moderately complex matter, and then learns that (a) one was subject to incredibly reliable brainwashing techniques designed to produce exactly the opinion one holds, and (b) one was dosed with special drugs which quite reliably both enhance people's susceptibility to brainwashing and leave them feeling mentally clear, and highly confident in their opinions. Should this discovery diminish one's confidence in one's philosophical view? It surely seems so. Yet any general principle which would mandate such a response would also seem to require, when the philosophical view in question is that very principle, that one lose

---

<sup>15</sup> The example is taken from Christensen (2009).

<sup>16</sup> See Frances (2010, 457). Frances's "Self-Application Objection" is essentially similar to what I'm calling the Inconsistency Argument.

<sup>17</sup> Elga (2010) also makes clear that the target of the Inconsistency Argument is not limited to theories of disagreement.

confidence in it. And this would get the Inconsistency Argument up and running. Thus the Inconsistency Argument would seem to rule out countless instances of epistemic modesty which are much less intuitively questionable than CV.

This surely provides a serious reason to be suspicious of the argument.<sup>18</sup> So maybe we shouldn't take the abstract possibility of self-undermining as a decisive objection to any of these principles. But before delving deeper into the question of what to make of the Inconsistency Argument, I'd like to look at a line of response suggested by Frances—a line which is intended to be independent of understanding why, or if, the Inconsistency Argument goes wrong.<sup>19</sup>

Frances begins by pointing out that there is a difference between believing a principle true and arranging one's doxastic life in accordance with the principle. With this in mind, he suggests the following course of epistemic action (adapted to CV): Withhold judgment on CV, but arrange one's doxastic life in accordance with it anyway. Now this might seem to amount to simply denying level-connection principles which would require one's beliefs about the correctness of epistemic rules to constrain one's doxastic practice. But that is not what Frances has in mind. For while he withholds judgment on whether his conciliatory principle is correct, he does believe that it's in the vicinity of the truth:

As already mentioned, I don't believe that [CV] is true. But I do think that it's closer to the truth than other principles (such as [a version of steadfastness]). I don't mean to suggest that it has a high degree of truth but less than 1 (where degree 1 is the full truth). It might be perfectly false, whatever that means. What I base my decision on, and have confidence in, is the idea that [CV] is in the *vicinity* of an important truth. I think that [CV] is a good rule of thumb, where all that means is that it's in the vicinity of a truth (endorsing a rule of thumb doesn't mean endorsing its truth). (2010, 459)

---

<sup>18</sup> For some more reasons, see Schechter (forthcoming).

<sup>19</sup> What follows is not exactly Frances's argument, which concerns a principle of his own that's a close relative of CV. In adapting his argument to the general case of CV, I think I'm remaining faithful to his intentions. For the original argument see Frances (2010, 457-9).

The idea, then, is not to dispense with level-connections. Rather, the idea is to endorse a claim about epistemic rules that is weaker than CV, but which is still strong enough to cohere with forming lower-level beliefs in a conciliatory way.

If this strategy worked, it would save something very much like CV from the difficulty the Inconsistency Argument poses. And in doing this, it would lessen the worry about the Inconsistency Argument engendered by its wide application, for we could presumably take a similar attitude toward the many extremely plausible principles that it seems to preclude. The strategy would provide a way for prescriptions of epistemic modesty and the Inconsistency Argument to coexist peacefully.

I don't think, though, that the strategy will work in the end. It does allow the person who's attracted to the arguments for CV to follow CV without flouting it by believing CV in the face of excellent epistemologists' disagreement. But the agent pursuing this strategy ends up running afoul of CV anyway. For she remains confident in the weaker proposition: that CV is in the vicinity of the truth—i.e., that CV is closer to the truth than steadfast views are. And this weaker claim is also denied by the excellent epistemologists who support steadfast views; they think that steadfast views are closer to the truth.

Moreover, although I'm quite confident that there are many excellent epistemologists working today who would deny that CV is closer to the truth than steadfast views are, it is worth noting that the difficulty with using the above strategy as a way of avoiding the Inconsistency Argument does not depend on this fact. For consider CCV—the view that CV is Close to the truth, and closer than steadfast views. Even if there weren't a sufficient number of actual excellent epistemologists who reject CCV, it's clear that there are possible situations in which there are. And these are situations in which CCV would entail that high confidence in CCV was irrational. And this, as we've seen, is all that's needed to launch the Inconsistency Argument. So the strategy of maintaining confidence in a weakened conciliatory view does not seem to me to provide a way of escaping the Inconsistency Argument.

If this is right, it does raise the pressure the Inconsistency Argument exerts on CV. However, it also reinforces our reason to think that there must be something wrong with the conclusion we seem to be driven to by the Inconsistency Argument, since it can't seem to get along with even weakened versions of the extremely plausible principles that, in some possible circumstances, call for their own rejection.

## 5. Partial Conciliation and the Self-Exemption Response

As noted above, Elga holds that CV is refuted by the Inconsistency Argument. But he does not think that this means that all highly conciliatory positions are refuted. In response to the Inconsistency Argument, Elga proposes a modified form of CV. According to this view, one should be conciliatory about almost all topics, but not about the correct way of responding to disagreement; this is what he calls the “partially conciliatory view”. If I were to adopt such a view, call it PCV, I might be very conciliatory with respect to my belief in Internalism. But I’d be dogmatic with respect to PCV itself. In fact, Elga suggests that in degree-of-belief terms, I should accord the correctness of the rule I adopt probability 1: I should be absolutely certain that it’s correct.

This, as Elga points out, may seem arbitrary at first. But he argues that it really isn’t. After all, the Inconsistency Argument applies to any epistemic policy that, in any circumstance, says that one should lose confidence in its correctness. As Elga puts it, “In order to be consistent, a fundamental policy, rule or method must be dogmatic with respect to its own correctness.”<sup>20</sup> So, since all acceptable fundamental rules are dogmatic with respect to their own correctness, it’s not ad hoc to take such a feature to be present in our view which tells us how to respond to disagreement.

I think there is something clearly right in Elga’s point that the justification for exempting PCV from its own scope is general, and thus that the charge of ad-hoc-ness is not clearly apt. Moreover, the suggestion that I remain fully confident in PCV, while being conciliatory about all manner of other things, allows me to respect the intuitions behind CV very widely, while also respecting the level-connecting idea that my doxastic practice should cohere with my higher-order views about rational belief formation.

Nevertheless, I think that there is something unsatisfying about the resulting position. And though I don’t want to press the charge of ad-hoc-ness, my reasons for dissatisfaction with PCV do stem from the kind of observations one might make in explaining why PCV seems ad hoc. In particular, it seems to me that the prescriptions of PCV in certain cases will be sharply

---

<sup>20</sup> Elga (2010, 183). Elga defines a “fundamental” method as “one whose application is not governed or evaluated by any other method”.



counterintuitive, and that these prescriptions will be counterintuitive for much the same reason that the prescriptions of strongly steadfast views are often counterintuitive. So even granting that there is a non-arbitrary reason for exempting PCV from its own scope, the view faces intuitive difficulties similar to those facing completely steadfast views.

Consider the sort of reasoning that might convince me that PCV was correct. It would include, in the first place, thinking through and evaluating the complex arguments and counterarguments offered for and against conciliationism in general. It would also include thinking through the meta-epistemological considerations adduced in mounting the Inconsistency Argument, and those adduced in support of adding the dogmatic self-exemption to deal with the problem that the Inconsistency Argument presents. All this thinking is, to all appearances, exactly the same type of thinking I do when consider whether to accept Internalism, or any other complex and controversial philosophical view. Clearly, this type of thinking is highly fallible. But if that's right, then it seems that I must take seriously the evidence I may get that I've made a mistake somewhere in my own thinking about PCV. And the disagreement of the many excellent epistemologists who reject PCV would seem to constitute just this sort of evidence.

The oddness of refusing to take this sort of evidence on board in the present case can be brought out by considering how remaining absolutely confident in PCV should fit into my general reflective view about myself. Suppose, that is, that I follow PCV and remain absolutely confident in its correctness, despite the fact that it's rejected by many epistemologists I respect, and even rate as my superiors in philosophical skill. How should I view my own reasoning on this topic? Should I think that while I'm generally only moderately reliable when I think about philosophy, nevertheless when I think about arguments for general conciliation, and for not being conciliatory about conciliation, I'm especially immune from error? That seems extremely dubious. There is nothing about this particular topic that would make my way of thinking about it special, or especially immune from my usual sort of blunders.

Should I count myself just lucky, then? This seems more natural: given my general fallibility in thinking philosophically, it would indeed be lucky if I, rather than all those more-talented philosophers who reject partial conciliation, am the one who is right this time. But can it possibly be rational for me to have absolute certainty that I'm the one who lucked out in this case? That, too, seems extremely unpalatable. On what basis could I conclude that I'm the one who got lucky, rather than those who reject PCV? Of course, if PCV is correct, then the direct

arguments on the topic actually do support PCV, and hence indirectly support the claim that I'm correct in this particular belief, and have not made a mistake. But that sort of support is available in any disagreement when I've in fact evaluated the basic evidence correctly; the intuitive appeal of conciliatory views of disagreement (and of other principles of epistemic modesty) flows from rejecting that sort of reasoning as begging the question.

Thus it doesn't seem to me that it would be rational for me to be highly confident (let alone certain) that I'm either very lucky or using especially reliable methods in thinking about the topic of rational responses to disagreement. And so PCV, despite fitting in a natural way with the Inconsistency Argument, does not seem to me to provide a satisfactory solution to our problem.

## **6. A Defense of Epistemic Modesty**

Let us take stock. The Inconsistency Argument poses a strong *prima facie* threat to CV. But it turns out that the problem is not that CV is intrinsically inconsistent: one could consistently obey CV in the face of disagreement by losing confidence in CV, but then continuing to follow it anyway. The problem is that doing this is inconsistent with plausible ways of taking beliefs in general to be rationally constrained by beliefs about what beliefs are rational. But even if no such level-connecting principle is entailed by CV, some such level-connection idea seems to be inseparable from the main motivation for CV. So there is a real tension inherent in conciliatory views of disagreement. Moreover, as we've seen, the tension extends to a myriad of other views that encode a certain kind of epistemic modesty: views that allow evidence that I've made an epistemic mistake in thinking about P to affect the degree of confidence it's rational for me to have in P. And we have seen that some initially attractive ways of reacting to the Inconsistency Argument do not fully succeed in dissolving this tension.

Of course, there may be other ways of dissolving the tension—perhaps with some more subtle level-connection principle that can motivate principles of epistemic modesty without enabling the Inconsistency Argument. But at this point, I can't see any.

One might, of course, give up entirely on epistemic modesty. But I think that such a radical approach would be misguided. We are fallible thinkers, and we know it. We know that it often happens that we evaluate the arguments and evidence on a certain topic—as carefully and

conscientiously as we possibly can—and reach the wrong conclusion. That is to say, we often make epistemic mistakes. And we know that simply looking over the arguments and evidence again, no matter how carefully and conscientiously, cannot be expected disclose our mistakes to us.

That being so, it seems clear that a person who was interested in having accurate beliefs, and, thus, in correcting her epistemic errors, would not be rational to let her confidence in P be unaffected by evidence that she was especially prone to making epistemic mistakes about P. It would be irrational even in instances where the person had in fact managed to avoid epistemic error in her original thinking about P. To give one example: suppose a doctor, after reaching a confident diagnosis based on a patient's symptoms and test-results, comes to realize that she's severely sleep-deprived, that she's under the influence of powerful judgment-distorting drugs, that she's emotionally involved with the patient in a way likely to warp her judgment, or that many sober clinicians, on the basis of the same symptoms and tests, have reached a contrary diagnosis. Perhaps she learns all of these things! In such a case, it seems quite clear to me that it would be highly irrational for her to maintain undiminished confidence in her diagnosis.<sup>21</sup> So I don't think that we may plausibly resolve the tension by denying epistemic modesty entirely.

One might, then, ask whether a less radical response is possible. Is there a way of defending the insights of CV, and other expressions of epistemic modesty, from the challenge posed by the Inconsistency Argument?

It seems to me that there is. It cannot be formulated precisely without having in mind precise forms of CV and other relevant principles. But the structure of the general idea can be illustrated using vague and rough approximations of these views. Here is a sketch:

First, we should recognize rational ideals of the following two types:

1. Respecting evidence of our epistemic errors

This sort of ideal requires, for example, that in typical cases where one is initially confident that P, and one encounters good evidence that one's initial level of confidence in P is higher than that supported by one's first-order evidence (say, for example, skillful thinkers who share one's first-order evidence about P are confident that not-P), one will give significant

---

<sup>21</sup> I have argued for this at greater length in (2010b).

credence to the claim that one's initial level of credence is too high. This sort of requirement applies even when one hasn't actually made an error in one's initial assessment of the evidence.

## 2. Level-Connection

This sort of ideal requires that one's confidence in P be constrained by one's beliefs about what level of confidence the evidence supports. For example, such an ideal may preclude being highly confident of P while simultaneously believing that that high degree of confidence is much higher than that supported by one's evidence.

Putting ideals of types 1 and 2 together will yield rational principles of epistemic modesty such as CV, Minimal Humility, and principles forbidding confident beliefs in many cases involving conclusions reached while sleep-deprived, or brainwashed, on certain kinds of powerful drugs, etc.

Next, we should recognize that the rational principles of modesty may apply to themselves. So in certain cases (for example, where one has strong evidence that one has made mistakes in thinking about such a principle), it may not be rational to have full confidence in its correctness. At that point, ideals of level-connection may exert rational pressure against fully obeying the principle of modesty. For insofar as one is rational to doubt the correctness of the principle of modesty, one may well be rational to believe, for example, that the level of confidence the principle has prescribed for some other proposition P is too low. In such cases, one must fail to respect one of the ideals fully. For following the level-connection ideal in this instance will mean raising one's confidence in P to a point higher than the principle of modesty would permit. And violating the principle of modesty will mean violating one of the ideals from which it flows. This is the problem exploited by the Inconsistency Argument.

But 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.

This position—call it the conflicting-ideals view—is not an entirely comfortable one. But I would argue that the discomfort it involves is not a new one—it arises quite independently of the Inconsistency Argument. In fact, it arises in a great many cases where an agent correctly

appreciates the import of her first-order evidence (for example, she sees that her evidence entails P, or that P is the best explanation for her evidence), but then receives powerful higher-order evidence that her take on the first-order evidence is mistaken. To the extent that she respects the higher-order evidence by reflecting it in her first-order beliefs (say, by lowering her confidence in P), her first-order belief will diverge from what is supported by the first-order evidence alone. In thus falling short of, e.g., respecting logic, or inference to the best explanation, her beliefs will fall short of certain rational ideals.<sup>22</sup> So the motivation for the conflicting-ideals view does not just come from wanting to avoid the Inconsistency Argument.

In the context of thinking about the Inconsistency Argument, the conflicting-ideals view offers significant attractions. Perhaps the most important one is that it allows us to avoid the absurdities entailed by blanket rejections of all expressions of epistemic modesty. These absurdities (such as the one illustrated in the doctor case above) should trouble even those who find highly conciliatory views of disagreement implausible.

Another attraction, it seems to me, is that seeing the Inconsistency Argument as involving conflicting ideals also avoids having to hold that certain propositions about difficult issues in epistemology are immune from rational doubt. One might question this—after all, won't an instance of the Inconsistency Argument show that, *at least for the most fundamental epistemic rules*, doubting their correctness will lead to inconsistency? I think it's worth pausing to examine this question.

Let us suppose, for argument's sake, that there is one absolutely comprehensive epistemic rule that encodes the epistemically best response to every possible situation; call it the Über-

---

<sup>22</sup> I have argued for the conflicting-ideals view at length in (2010a), and in more compressed forms in (2007) and (2011). I also defended a particular version of a level-connection principle this way in (2010b), but the unpublished work by Elga referred to above has persuaded me that this last invocation of the strategy may have been unnecessary. Joshua Schechter (forthcoming) writes that he suspects that this sort of diagnosis applies to what he calls the “fixed-point argument”—a version of what I’m calling the Inconsistency Argument. But he expresses some reservation, since he sees it as a “radical view”. But if the view is in fact motivated independently of the Inconsistency Argument, invoking it here doesn't require making any new radical commitments. And as we've seen, other natural treatments of the Inconsistency Argument involve more radical departures from intuitive judgments about epistemic rationality.

rule.<sup>23</sup> The Über-rule will surely qualify as a fundamental rule, in the sense that its application is not governed by any other rule: by construction, what the rule says an agent is most rational to believe in a certain situation is exactly what the agent is most rational to believe in that situation, so an agent cannot rationally diverge from the rule's prescriptions. Such a rule might not be stateable in any compact natural way; it might well be enormously complex, and exceedingly difficult to formulate and think about. But if anyone were lucky enough to formulate it correctly, it's a safe bet that it would be controversial among skilled epistemologists. And given the difficulty of the topic, and the attendant controversy, it intuitively would seem irrational for the formulator to be extremely confident that it was exactly correct. So it's very plausible that the epistemically best option for the agent—*i.e.*, *the option recommended by the Über-rule*—would involve the agent having less than full confidence that the rule she's formulated is correct.

However, once we allow this, one might worry that trouble will ensue. If the agent continues to follow the Über-rule while doubting its correctness, it seems inevitable that she will in some cases violate the sort of level-connection ideal we've been discussing. So to the extent that the agent rationally doubts that the Über-rule is correct, it seems that she will be rationally required to violate a level-connection ideal. Does this mean that our initial assumption—that there are cases where an agent should doubt the correctness of the Über-rule—must be rejected, so that we must after all hold the Über-rule immune from rational doubt? Not necessarily. For on the conflicting-ideals view, it may be that the epistemically best option for the agent will involve violating a (perfectly legitimate) level-connection ideal. There is a sense in which this violation is epistemically regrettable, but that doesn't mean that there was a better option for the agent. After all, having full confidence in the correctness of the Über-rule would fly in the face of powerful evidence (such as the disagreement of excellent epistemologists) bearing on her own fallible thinking about abstract issues in epistemology. Disrespecting that evidence would also involve violation of an epistemic ideal.<sup>24</sup>

---

<sup>23</sup> I do not wish to argue that there is, in fact, an Über-rule which specifies the unique best response to every situation. One might hold that rationality is permissive, and that more than one response will be tied for best in certain situations. One might even hold that in certain situations, several permissible responses will be incomparable with one another. I'm raising the possibility of the Über-rule to bring out sharply a certain possible problem with the conflicting ideals view.

<sup>24</sup> Compare the discussion of Über-rules in Aarnio (forthcoming).

I should note that the conflicting-ideals view does not entail that, in every particular case, the option of violating a level-connection ideal will be better than the option of disregarding evidence of one's error. The view leaves room for both types of response. The point is just that the view allows for the possibility (which strikes me as very plausible) that the most rational option facing an agent will sometimes involve having some degree of doubt about the correctness of even the most fundamental epistemic rule. So we are not stuck with having to say that the epistemically best option for all agents must include having absolute confidence in the correctness of a particular subtle and complex position in epistemology.

Finally, the conflicting-ideals view explains how defenders of CV can hold on to the motivation for their views, while acknowledging that the level-connecting ideals that lie behind their views must, in certain circumstances, be violated. For the conflicting-ideals view recognizes that an ideal that must sometimes be violated may yet have force. If we ask *why* the disagreement of other competent thinkers with the same evidence should affect my confidence, the correct explanation may still be that since their disagreement is evidence that my initial belief was based on an epistemic error, it creates rational pressure to give credence to the claim that my initial belief was based on error, and that (as ideals of level-connection would have it) this creates rational pressure to back off of that initial belief to at least some extent.

It's worth emphasizing that the advantages I've been advertising are not just available to defenders of conciliatory views of disagreement. For one thing, they are equally available to defenders of even moderately steadfast views of disagreement. Such views may see first-order considerations as being more robust in the face of the higher-order doubts prompted by disagreement, but to the extent that they recognize that the disagreement of others (even, e.g., large numbers of demonstrably smarter agents) should temper one's confidence to some extent, they depend on level-connecting principles for their motivation. And, as we've seen, such views have just as much trouble with the Inconsistency Argument as does CV. And the same point also goes, of course, for non-disagreement-based prescriptions of epistemic modesty, for example, in the face of evidence that one has been brainwashed or drugged. So the conflicting-ideals view is not just good for conciliationists.

There are, of course, concerns one might reasonably have about adopting this approach. One such worry is that it would make it much harder to give a theoretically tractable account of

rational belief.<sup>25</sup> For example, consider an account on which rational credences are given by some sort of epistemic probabilities. This sort of account seems to offer important insight into the structure of rational belief—it shows how credences are constrained by logic. And it does this by means of a clear, precisely-describable formal condition. But one might think that the conflicting-ideals view poses severe difficulties for this picture. After all, it would seem that a rational agent's confidence in some logical truth—which, if one's credences are probabilistically coherent must be 1—may well be rationally undermined by evidence of her epistemic error. Perhaps she gets told by an expert logician that it's not a logical truth, or learns she was drugged while thinking through the proof, etc. The conflicting-ideals view may allow (and I think it should allow) that in some such cases, the rationally best credence for the agent will fall short of 1. In such cases, the rationally best credences will not be probabilistically coherent.

So we do lose the ability to take probabilistic coherence as a necessary condition on the beliefs it would be most rational for an agent to adopt. Moreover, it seems doubtful that some other relatively simple formal condition—or even some precisely stateable informal condition—will be able to take its place. Finally, we have no reason to suppose that, if an ideal of probabilistic coherence is balanced against other ideals, say by means of an Über-rule, that the recipe for combining them will be capturable in any tidy formula.

I think that this is an important point. I suspect that insofar as one's account of epistemic rationality takes agents' reflection on their own fallibility seriously—in the sense of encoding epistemic modesty—it is unlikely to provide a clear, relatively simple account of rational belief. How serious a problem is this?

I cannot enter here into a full discussion of the methodological issues raised by this worry. But it seems to me that there are a couple of points to keep in mind.

The first is that we don't lose the ability to theorize about the ideals that contribute to epistemic rationality; and these ideals may well include clear, precisely-describable formal conditions. For example, we're free to put aside worries about self-doubt in thinking about how logic constrains rational credence. Such thinking may convince us that a formal condition such

---

<sup>25</sup> Thanks to Timothy Williamson for prompting me to address this worry. See Aarnio (forthcoming) for an extended development of a related worry about any way of allowing for epistemic modesty (at least in the context of an epistemology that sees rational belief in terms of applying correct epistemic rules).



as probabilistic coherence correctly encodes the rational pressure that logic puts on degrees of belief.<sup>26</sup>

In fact, the conflicting-ideals view is in a way particularly hospitable to theorizing about rationality in a way that utilizes a condition such as probabilistic coherence. A major obstacle to this sort of theorizing comes from considering cases where agents seem rationally required to violate coherence. But the conflicting-ideals view allows us to see how the ideal can have force even in such situations, by allowing that there's something epistemically imperfect about the beliefs of the agent who violates coherence as a result of rationally accommodating evidence about her own malfunction.

So it is important to see that adopting the conflicting-ideals view does not amount to giving up on epistemic theorizing. And it does not render useless the formal tools many have employed to illuminate aspects of epistemic rationality; in fact, it may help make room for just this sort of theorizing.

Another point to notice is that, insofar as we're persuaded that in cases such as the doctor case discussed above, there would be something epistemically defective about the doctor's undiminished confidence in her diagnosis after she receives the powerful evidence of her own malfunction, it seems that epistemologists should be interested in explaining what that epistemic defect is. Of course, one could respond by walling off a certain aspect or dimension of epistemic appraisal—for example, by deciding to theorize about a notion of rationality on which the rationality of an agent's confidence in P was by definition not affected by way of evidence raising doubts about whether the agent made a mistake in thinking about P. One might thus develop a sanitized notion of rational belief, on which the doctor's undiminished high confidence in her diagnosis, in the face of the strong evidence of her cognitive malfunction, was perfectly rational. One would then need to say that the doctor's undiminished confidence embodied some other epistemic defect. But one would still be left with the task of theorizing about that defect. So it seems that at some point, epistemology must confront the problem raised by this sort of evidence. And to my mind, the intuitive irrationality of the doctor's undiminished confidence in

---

<sup>26</sup> For more on this point, see Christensen (2007).

this situation speaks strongly in favor of doing that theorizing in the context of describing rational belief.<sup>27</sup>

Another worry one might have about the conflicting-ideals view is that it's too promiscuous, like a general-purpose "get out of jail free" card for perpetrators of crimes against coherence. I think that this is also a reasonable concern. But the concern may perhaps be mitigated by noting that the cases in which ideals conflict share an important feature: they all involve the results of agents reflecting critically on their own thinking. Perhaps it is not so surprising that insofar as it is rational to take seriously one's critical assessments of one's own beliefs, certain kinds of incoherence will result.

Suppose a certain rational ideal applies to beliefs. Suppose also that one can't always rationally be certain whether the ideal is in fact a correct one, or, alternatively, whether one is actually obeying this ideal. In particular, suppose that it can be rational to have such doubts, even in cases where the ideal is correct, and one is in fact obeying it. Suppose also that this sort of critical reflection on one's own beliefs is not merely an idle exercise—that rational doubts about whether one's beliefs meet the correct ideals may make it rational to change the beliefs in question. In such cases, one may well come under rational pressure to violate a rational ideal.

On this picture, allowing for the possibility of conflicting ideals reveals no penchant for promiscuity. The idea is not that everything is permitted—in certain cases it will be positively irrational to satisfy a particular ideal. The idea is just to make room for modesty. The conflicting-ideals view simply allows us to recognize the rationality of acknowledging, and then taking serious account of, the possibility that we've fallen short of epistemic perfection. If we can accommodate that sort of modesty in our account of rational belief, it seems to me that it will be well worth the price of abandoning the hope that some cleanly specifiable notion of coherence is satisfied by the maximally rational response to every evidential situation.

## References

---

<sup>27</sup> Compare the last section of Aarnio (forthcoming).

- Aarnio, M. L. (forthcoming), "The Limits of Defeat," *Philosophy and Phenomenological Research*.
- Christensen, D. (2011), "Disagreement, Question-Begging and Epistemic Self-Criticism," *Philosophers' Imprint* 11, no. 6.
- . (2010a), "Higher-Order Evidence," *Philosophy and Phenomenological Research* 81.1: 185-215.
- . (2010b), "Rational Reflection," *Philosophical Perspectives* 24: 121-140.
- . (2009), "Disagreement as Evidence: The Epistemology of Controversy", *Philosophy Compass* 4: pp. 756-767.
- . (2007), "Does Murphy's Law Apply in Epistemology? Self-Doubt and Rational Ideals", *Oxford Studies in Epistemology* 2: 3-31.
- Elga, A. (2007), "Reflection and Disagreement," *Noûs* 41: 478-502.
- . (2010), "How to Disagree about how to Disagree," in Feldman, R. and T. Warfield.
- Feldman, R. and T. Warfield, eds. (2010), *Disagreement* (New York: Oxford University Press).
- Field, H. (2000), "Apriority as an Evaluative Notion," in P. Boghossian and C. Peacocke, eds., *New Essays on the A Priori* (New York: Oxford).
- Frances, B. (2010), "The Reflective Epistemic Renegade," *Philosophy and Phenomenological Research* LXXXI, 2: 419-463.
- Kelly, T. (2005), "The Epistemic Significance of Disagreement," *Oxford Studies in Epistemology* 1: 167-196.
- . (2010), "Peer Disagreement and Higher-Order Evidence," in Feldman and Warfield.
- Lackey, J. (2010), "A Justificationist View of Disagreement's Epistemic Significance," in A. Haddock, A. Millar and D. Pritchard, eds., *Social Epistemology* (Oxford: Oxford University Press): 298-325.
- . (2010a), "What Should We Do when We Disagree?", *Oxford Studies in Epistemology* 3: 274-29.
- Lewis, D. (1971). "Immodest Inductive Methods," *Philosophy of Science* 38: 54-63.
- Schechter, J. (forthcoming), "Rational Self-Doubt and the Failure of Closure," *Philosophical Studies*.
- Sosa, E. (2010), "The Epistemology of Disagreement," in A. Haddock, A. Millar and D. Pritchard, eds., *Social Epistemology* (Oxford: Oxford University Press).

Weatherson, B. (2007): “Disagreeing about Disagreement,” <http://brian.weatherson.org/DaD.pdf>,  
incorporated into his contribution to this volume.

Wedgwood, R. (2010) “The Moral Evil Demons,” in Feldman and Warfield.

Weiner, M. (2007): <http://mattweiner.net/blog/archives/000781.html>.