

Disagreement and the Burdens of Judgment

Thomas Kelly
Princeton University

1. Some cases

Case 1: Intrapersonal Conflict. Suppose that you suddenly realize that two beliefs that you hold about some subject are inconsistent with one another. Prior to becoming aware of the conflict, you were quite confident of each. Indeed, let's suppose that you were more or less *equally* confident that they were true. Now that you are aware of the conflict, how should you revise your beliefs?

A possible answer: in any case of the relevant kind, you are rationally required to abandon both beliefs until you acquire further evidence. In particular, it would be unreasonable to retain one of the two beliefs while abandoning the other.

A better view: In *some* cases of intrapersonal conflict, the reasonable thing to do might be to abandon both beliefs until further evidence is acquired. But in other cases, it might be perfectly reasonable to resolve the conflict by dropping one of the two beliefs and retaining the other. What would be a case of the latter kind? Paradigmatically, a case in which one of the two beliefs is well-supported by your total evidence but the other is not. If your total evidence strongly supported one of the two beliefs before you became aware of the conflict, then it might very well be reasonable to retain that belief even after you realize that it is inconsistent with something else that you have confidently believed up until now.

Case 2: Disagreeing with A Great Dead Philosopher. You regularly teach a certain classic philosophical text. In the course of doing so, you have come to have a high opinion of the author's judgment. It's not simply that you think that he manifests genuine intellectual virtues such as creativity and imagination in the course of pursuing the questions with which he is concerned, but that his judgment is generally quite reliable.¹

¹Consider the way in which one might genuinely admire the ingenuity displayed by

Indeed, you would not claim to have better judgment about the domain than he does. Nevertheless, against this background of genuine respect, you believe that his discussion of some particular issue is uncharacteristically weak. It's not that you think that you have uncovered a clear fallacy, or something akin to a clear fallacy, the kind of thing such that if you could only travel back in time and present it to the author, a modicum of intellectual humility would compel him to change his mind. Rather (what I take to be a much more typical case), you simply think that the author has overestimated the force of his own arguments and underestimated the force of the objections that he considers. You thus conclude that the author's case for his conclusion is quite weak. Could this attitude be reasonable?

A possible answer: No, it couldn't. After all, the author himself evidently took his case to be quite compelling, and you make no claim that your judgment is generally superior to his in this domain. Because of this, you should treat his judgment that his case is compelling, and your initial judgment that his case is quite weak, evenhandedly. Thus, you should be more or less agnostic about the merits of the author's case.

A better view: It might very well be reasonable for you to be confident that the author's case is not compelling. What would be circumstances in which this is a reasonable attitude on your part? Suppose that the author *really has* overestimated the extent to which his arguments tell in favor of his conclusion and underestimated the force of the objections that he considers, and that this is something that you've correctly picked up on. In those circumstances, you might very well be justified in having a low opinion of the author's discussion, notwithstanding his contrary opinion.

Case 3: Disagreeing with One's Past Self. Like many other philosophers, I sometimes write relatively detailed notes in the margins when reading a philosophical text. More eccentrically, I've also long maintained the practice of dating the cover page of every book that I read, each time that I read it. (Thus, among the not especially interesting autobiographical facts to which I currently have access is that I first read Descartes' *Meditations* in March of 1991.) Because of these practices, I'm often in a

Leibniz or Spinoza in their metaphysical writings while nevertheless thinking that they are quite unreliable about matters of ontology. In my own case, the attitude that I take towards Frege's work on the epistemology of mathematics would provide a much better example of the kind of thing that I have in mind.

position, upon rereading a philosophical text, to compare my current view of some author's arguments with my past view of those arguments. Perhaps unsurprisingly, I often find that I agree with my past self, even when the past self in question is a relatively distant one. For example, I don't think particularly highly of Descartes' arguments for the existence of God now, and I'm pleased to report that it seems that my eighteen year old self didn't think too highly of them either.

But on occasion, my past self and I disagree. Of course, it's easy enough for me to write off the opinions of my eighteen year old self as the misguided ruminations of a philosophical novice. On the other hand, and as much as I'd like to believe otherwise, I don't believe that my philosophical judgment now is appreciably better than it was, say, five years ago, back when I was a lowly assistant professor. Rather, I suspect that at this point I've pretty much leveled off when it comes to the amount of sophistication that I bring to bear when I critically engage with a text. Consider then those cases in which I disagree with *my philosophically mature self*: upon re-reading some philosophical text, I find that I'm inclined to change my mind about an author's discussion, which now seems significantly weaker (or stronger) than it did in the past.

Suppose that I go ahead and change my mind. Perhaps my awareness that I once thought differently tempers my current confidence a bit, since it reminds me of my fallibility with respect to judgments of the relevant sort. But the opinion that I settle on now is significantly closer to what I would have thought if I had simply made up my mind in complete ignorance of what I used to think, than it is to the opinion that I used to hold. Question: could this be reasonable on my part?

A possible answer: No, it would never be reasonable for me to do this. After all, I don't think that my judgment has significantly improved when it comes to this general kind of question; nor do I claim to possess any 'silver bullet' piece of evidence that I lacked then. Given this, it would be completely arbitrary to privilege the opinion of my current self over the opinion of my past self. Therefore, I should be even-handed and give equal weight to the opinion of my philosophically mature past self and the opinion that I'm now inclined to hold.

A better view: in at least some cases, it might very well be reasonable for me to change my mind significantly, and adopt an opinion that is relatively close to the view

that I would hold if I ignored my past opinion entirely, and relatively far from the view taken by my past self. What would be a case of this kind? Suppose that it's not simply that it *seems* to my current self that my past self gave too little (or too much) weight to the author's arguments, but that it's *true* that my past self did this. That is, suppose it's a case in which my current self is right to think that my past self misjudged the probative force of the evidence and arguments that are presented in the text. In cases of this kind, it doesn't follow that I'm unreasonable, even if I don't treat my past opinion and the opinion that I'm inclined to hold even-handedly.

But how can I be sure that I'm getting things right now, as opposed to then? There is now way for me to be sure of this. If I favor the assessment of my current self, but my past assessment was more accurate, then the opinion that I end up holding will be unreasonable (or at least, less reasonable than another opinion that I might have held instead). But that's simply the familiar fate of someone who misjudges his or her evidence.

I'm interested in what to say about all of these cases, and what others think about them. But I put them on the table primarily as a warm-up for thinking about the case with which I'll be primarily concerned for the rest of the paper.

Case 4: Peer Disagreement. Suppose that you and I have been exposed to the same evidence and arguments that bear on some proposition: there is no relevant consideration that is available to you but not to me, or vice versa. For the sake of concreteness, we might picture the following:

You and I are attentive members of a jury charged with determining whether the accused is guilty. The prosecution, following the defense, has just rested its case.

Suppose further that neither of us has any particular reason to think that he or she enjoys some advantage over the other when it comes to assessing considerations of the relevant kind, or that he or she is more or less reliable about the relevant domain. Indeed, let's suppose that we possess significant evidence that suggests we are likely to be more or less equally reliable when it comes to questions of the relevant kind. Because we're aware of this, if we had been asked in advance of the trial which one of us is more likely to be wrong in the event of a disagreement, we would have agreed that we were equally

likely to be wrong.² Nevertheless, despite being (apparent) peers in these respects, you and I arrive at different views about the question on the basis of our common evidence. For example, perhaps I find myself quite confident that the accused is guilty while you find yourself equally confident that he is innocent.

Suppose next that, upon learning that I think that the accused is guilty, you reduce your confidence in his innocence. However, even after you take my opinion into account, it still seems to you that on balance the evidence suggests that he is innocent. You still regard it as significantly more likely that he is innocent than that he is guilty, to the point that you can correctly be described as retaining your belief in his innocence. Question: in these circumstances, is there any possibility that this is a reasonable response on your part?

*A possible answer: **The Equal Weight View/Conciliationism*** (Cf. Elga 2007; Christensen 2007a and forthcoming; Feldman 2003, 2006, 2007; Kornblith 2010; Cohen [ADD to this list?]). No, there isn't. In any case of the relevant kind, you are rationally required to abandon your original belief and retreat to a state of agnosticism. (I'm required to do the same.) Given the relevant symmetries, you should give *equal weight* to my view as to yours; thus, given that initially I'm confident that the accused is guilty while you're equally confident that he is not, the uniquely rational stance is for us to suspend judgment about the issue. As Richard Feldman puts it:

[C]onsider those cases in which the reasonable thing to think is that another person, every bit as sensible, serious, and careful as oneself, has reviewed the same information as oneself and has come to a contrary conclusion to one's own... An honest description of the situation acknowledges its symmetry... In those cases, I think, the skeptical conclusion is the reasonable one: it is not the case that both points of view are reasonable, and it is not the case that one's own point of view is somehow privileged. Rather, suspension of judgment is called for (2006: 235)

*A better answer: **The Total Evidence View*** (Kelly 2007). Yes, there is. Whether it's reasonable for you to believe that the suspect is innocent even after learning that I think

²Cf. Elga's (2007) account of what it is to treat someone as an *epistemic peer* in his sense.

otherwise is not something that can be determined, given only the facts about the fiction provided. What are some circumstances in which your belief might be reasonable? Suppose that the original evidence with which we are presented strongly supports the view that the suspect is innocent. Your original belief is a rational response to what was then our total evidence; mine was not. (Against a general background of competence, I commit a performance error.) After you learn that I think that the accused is guilty, your total evidence has changed: it's now on the whole less supportive of the view that he is innocent than it was previously. It's thus reasonable for you to reduce your confidence to at least some degree. Still, the total evidence available to you then might very well make it more likely that the suspect is innocent than that he is guilty, to the point that it's reasonable for you to believe that he is guilty. In any case, there is certainly no guarantee that the uniquely reasonable response on your part is to retreat to a state of agnosticism between your original opinion and my original opinion, as the Conciliationist suggests.

2. Conciliationism

Conciliationism plays a central role in structuring the epistemology of disagreement literature. Some prominent contributors to the literature endorse the view, but even those who reject it often use it as the background against which to develop their preferred alternatives. If any view deserves the title of the View to Beat, it is this one. However, despite its prominence in the literature, and although the animating intuition seems straightforward enough, it's not clear exactly what the view actually *says*. Among other things, we currently lack an adequate precisification of what it means to give 'equal weight' to someone else's opinion and to one's own.³ David Christensen, who in my

³ On this point, see especially Jehle and Fitelson, "What is the 'Equal Weight View'?" In several places in Kelly (2007), I slipped into interpreting 'giving equal weight to your peer's opinion' as a matter of *averaging* your initial credence and your peer's initial credence in order to arrive at an updated credence. As a number of people have pointed out, however, this "arithmetical mean" interpretation is not plausible as a general interpretation. It works well enough in the special case with which I was primarily concerned, in which you invest a certain credence in *p* and I invest the same credence in not-*p*. (In that case, averaging the original credences leads us to converge on a revised credence of .5, representing perfect agnosticism, which is where proponents of the view think we should be.) However, the interpretation breaks down in cases where we

estimate has done at least as much as anyone else to develop this general approach to disagreement, freely admits that there is as of yet no suitably general and determinate principle of belief revision on the table (forthcoming, p.x). For this reason, I want to be quite clear about what *I* will mean by ‘Conciliationism’. For purposes of this paper, a view counts as Conciliationist if and only if it entails that suspending judgment is a *necessary* condition for being reasonable in a canonical case of peer disagreement, i.e., any case that has the same structural features as Case 4 above. (In a framework employing degrees of belief, the necessary condition should be interpreted so as to require a stance of agnosticism, i.e., a degree of belief of approximately .5). Notably, a Conciliationist need not hold that suspending judgment in such circumstances is also a *sufficient* condition for reasonableness, although some Conciliationists might very well hold that it is.⁴

are not centered around the midpoint in this way. For example, suppose that I perform a calculation and invest 90% confidence in the number at which I arrive (I give some weight to the possibility that I’ve made a mistake.) I then learn that you arrived at the same number independently, and currently invest 80% confidence in that answer. Intuitively, learning that you arrived at the same number should make me *more* confident of my answer, as opposed to less confident, which would be the outcome of averaging our credences [REF to DC on this point]. And surely someone who recommends that I give equal weight to both of our opinions should not be understood as saying otherwise. Thus, the arithmetical mean interpretation is not viable across the board. (I don’t believe, however, that this point detracts from whatever force my objections in the aforementioned paper possess.)

⁴ This point is emphasized by Christensen (forthcoming), in the course of developing a version of Conciliationism designed to avoid some of the objections offered in Kelly (2007). Christensen suggests that the Conciliationist should say that even if one responds to the disagreement by adopting the rationally required attitude, that attitude might still be unreasonable, if one’s initial belief was unreasonable given one’s initial evidence. For criticism of this way of understanding the view, see Cohen [this volume] and Kelly (manuscript).

A further note about terminology is in order here. Occasionally, “Conciliationism” is used in an extremely inclusive way, so that any view according to which one must give at least some (even extremely minimal) weight to the opinion of a peer in a canonical case of peer disagreement counts as a species of Conciliationism. (See, e.g., Elga (2010).) On this inclusive usage, the Total Evidence View, no less than the Equal Weight View, counts as a species of Conciliationism. As the above makes clear, however, I will use the term in a much less inclusive way, following the practice of Christensen’s (2009) helpful survey of the literature.

Elsewhere I've argued against Conciliationism and defended The Total Evidence View at some length. Here I want to attempt to get a bit deeper, in a way that builds on the insights of some other contributors to the debate. Although I am perhaps prone to bias when it comes to evaluating the health of the disagreement literature, I do think that there has been some discernible progress. It's not simply that there has been some convergence among prominent representatives of competing approaches with respect to the ultimate issues. It's also that, with respect to the major divisions that persist, there is greater clarity about what underwrites these divisions, and what it would take to resolve them. An example of the latter: in his recent critical survey of the literature (2009), Christensen identifies a principle, 'Independence', that he thinks underlies the division between those who accept Conciliationism and those who reject it:

Independence: In evaluating the epistemic credentials of another person's belief about P, in order to determine how (if at all) to modify one's own belief about p, one should do so in a way that is independent of the reasoning behind one's own initial belief about p (2009:758).

According to Christensen, the dispute between Conciliationists and Non-conciliationists is explained by the fact that the former accept, while the latter reject, a principle of this sort.⁵ Although he is a Conciliationist and I am not, I find Christensen's diagnosis of the debate relatively plausible. Thus, I'd like to consider this principle at some length and with some care.

3. Independence

Notice that Independence is a relatively general epistemic principle, one which says nothing about the case of peer disagreement in particular. According to Christensen

Conciliationism will result from combining this sort of principle with the thought that, to the extent that one's dispute-independent evaluation gives one strong reason to think that the other person is equally likely to have evaluated the evidence

⁵ Christensen himself endorses Independence, and, as he notes, extremely similar principles have been explicitly endorsed by other Conciliationists, e.g., Elga (2007) and Kornblith (2010).

correctly, one should (in the case where one is quite confident that p , and the other person is equally confident that not- p) suspend belief (or adopt a credence close to .5) in p (2009: 758-759)

But why should we accept anything like this principle? Here is Christensen, one more time:

The motivation behind the principle is obvious: it's intended to prevent blatantly question-begging dismissals of the evidence provided by the disagreement of the others. It attempts to capture what would be wrong with a P-believer saying, e.g., "Well, so and so disagrees with me about p . But since P is true, she's wrong about p . So however reliable she may generally be, I needn't take her disagreement about p as any reason at all to change my belief".

There is clearly something worrisome about this sort of response to the disagreement of others. Used as a general tactic, it would seem to allow a non-expert to dismiss even the disagreement of large numbers of those he took to be experts in the field ($p.x$).

In several places, Christensen employs his own paradigm example of a disagreement between peers, "The Ordinary Restaurant Case", in order to illustrate how we should apply Independence.⁶ In the Ordinary Restaurant Case, you and I independently calculate our shares of the dinner tab (we've agreed to divide the check evenly among everyone who was at dinner). We know, on the basis of substantial track record evidence, that we're more or less equally competent when it comes to performing this general kind of calculation (in our long history of dining together, we almost always come up with the same number, but on those occasions when we've come up with different numbers, each of us has turned out to be the one who was correct approximately half the time). On this occasion, you arrive at the number \$43 while I arrive at the number \$45. A widely shared intuition is that, upon discovering this, both of us (including the person who in fact reasoned impeccably, assuming that one of us did) should become much less confident of his or her original answer, and that indeed, each of us should divide our credence between the two answers more or less equally. Notice how, applied to this case, Independence yields the widely shared intuition. For Independence instructs each of us to set aside the reasoning which led us to our original answer in evaluating the "epistemic

⁶ The Ordinary Restaurant Case was first introduced in Christensen (2007a). It also appears in Christensen (2009, 2010, forthcoming).

credentials” of the other person’s belief; once we’ve done this, we’re left with the knowledge that the other person is, in general, more or less equally reliable when it comes to this kind of calculation. And this in turn suggests that each of us should treat the two original opinions even-handedly in arriving at a new view. (At least, until we perform the calculation again, or consult a calculator.) Elsewhere, Christensen writes of the need to “bracket” the reasons and evidence on the basis of which one reaches one’s original answer, once one becomes aware of the disagreement.⁷ For notice that, if in fact you reasoned impeccably in arriving at your original answer, then the facts from which you reasoned (that the total bill is n dollars; that m people have agreed to divide the check evenly, etc.) literally *entail* the correct answer. So if such facts are among the evidence you have to go on in evaluating my belief, then they would seem to provide a basis for discounting my opinion entirely. But according to Independence, I should set aside such facts when evaluating your belief.

What should we make of Independence? First, a couple of preliminary remarks. How should we understand talk of “evaluating the epistemic credentials of another person’s belief about p ?” An obvious first thought is that such evaluation is a matter of judging the epistemic status of the person’s belief: for example, making a judgment about how reasonable that belief is. But on reflection, it’s clear that “evaluating the epistemic credentials of another person’s belief” will have to include considerably more than mere judgments of reasonableness, given the role that such evaluation is supposed to play in guiding revision of one’s own beliefs. For on anyone’s view, the mere fact that one evaluates someone else’s opinion as perfectly reasonable completely leaves open how much weight (if any) one should give to that opinion. For example, suppose that my loathing of the butler leads me to frame him for some crime that he didn’t commit. Suppose further that I execute my plan impeccably: due to my efforts, the authorities and many members of the general public come to possess large quantities of misleading evidence, all of which suggests that the butler committed the crime. On the basis of this evidence, you become extremely confident that the butler did it. When I subsequently meet you and note with satisfaction how confident you are of the butler’s guilt, I might very well judge that your belief is perfectly reasonable. Nevertheless, in these

⁷ On “bracketing”, see especially his (2010).

circumstances, the mere fact that you reasonably believe that the butler committed the crime is no reason at all for me to be even slightly less confident of my belief that the he did not commit the crime. On the other hand, it's also true, on anyone's view, that when I encounter a person who I take to be both *better* informed than I am about some question and perfectly reasonable in believing as she does, the fact that she believes as she does gives me a reason to revise my view in her direction.⁸

The moral: given that “evaluating the epistemic credentials of another's belief that p” is supposed to play a role in potentially guiding one's own belief revision, such evaluation will have to go considerably beyond judgments about whether the other person is reasonable in believing as she does. Such evaluation will also require judgments about the quality of her evidence and how well-informed she is. This point will be important for us later on; for now, I want to flag it and move on.

A second preliminary point concerns a worry about the way in which Christensen formulates Independence. For there is an aspect of that formulation that threatens to severely limit its applicability (or at least, an aspect that makes it unclear how the principle should be applied in a significant range of cases). Independence requires that, when one evaluates another's belief that p, one bracket “the reasoning behind one's initial belief about p”. Talk of “the reasoning behind one's initial belief” is easiest to understand in cases that closely resemble Christensen's favorite example, the Ordinary Restaurant Case. In that case, there really is some identifiable, relatively discrete piece of reasoning that leads one to a particular belief. But many beliefs, including many of the beliefs that philosophers are ultimately concerned with in the disagreement literature (e.g., the kinds of extremely controversial beliefs that people hold about history, politics, religion and philosophy) are often not easily understood as the output of some discrete process of reasoning.

Consider, for example, two different bases for atheism. Undoubtedly, some atheists believe as they do because of some relatively discrete piece of reasoning. (We might

⁸ If I know that she's both better informed than I am and perfectly reasonable in believing as she does, shouldn't I simply adopt her opinion as my own? Not necessarily, for her being better informed than I am is consistent with my having some relevant evidence that she lacks, and this can make it reasonable for me not to simply adopt her opinion.

think here of someone who arrives at the view that there is no such being as God by reasoning from some relatively small set of premises, at least one of which refers to the existence of evil in the world.) But alternatively, one might disbelieve in God for the following reason: given everything else that one takes to be true about reality, one judges that it's extremely improbable that any such being exists. On Christensen's formulation, it's much easier to see how Independence applies to the former sort of atheist than to the latter. (Presumably, in evaluating someone else's belief about God, the atheist is not supposed to bracket *everything* that she takes to be true about reality, even if the reason she invests extremely low credence in the proposition that God exists is the fact that that proposition has extremely low probability given everything else she believes). Perhaps it still makes sense to talk about "the reasoning behind the belief that God does not exist" in the second case as well as the first. But if so, we should not underestimate the difference between the two cases.⁹

Having raised this issue, I will for the most part ignore it in what follows. But I do want to insist on a statement of Independence that has clear application to a case in which one's initial belief is based on one's assessment of a given body of evidence or information. For this purpose, I propose the following, which I take to be completely in the spirit of Christensen's original statement:

Independence*: In evaluating the epistemic credentials of another person's belief about P, in order to determine how (if at all) to modify one's own belief about P, one should do so in a way that is independent of one's assessment of those considerations that lead one to initially believe as one does about p.

Although I still have some worries of detail about this formulation, I'll waive these in what follows.

The first substantive point that we should note about Independence* (or for that matter, Independence) is that it is an extremely strong principle. Suppose that I possess a great deal of evidence¹⁰ that bears on the question of whether the Holocaust occurred; I look it over and judge correctly that this body of evidence strongly confirms that the

⁹ NOTE THAT DC himself is well aware of this issue.

¹⁰ Here, we can think of my evidence as a body of information, or a collection of facts, to which I have cognitive access.

Holocaust occurred; on the basis of that assessment, I invest a correspondingly high amount of credence in the proposition. I then encounter a Holocaust denier. (For purposes of the example, let's imagine that this person is quite reliable when it comes to matters that are unrelated to the Holocaust.) In evaluating the epistemic credentials of his belief that the Holocaust never occurred, Independence* would have me bracket my assessment of all of those considerations which led me to believe that the Holocaust did occur. An obvious question is whether, once I do this, I'll still have enough left to go on to offer an evaluation of the Holocaust denier's belief. A second question is why, even if I do have enough left to go on to arrive at an evaluation, we should think that the evaluation that I come up with under those conditions is worth anything.

Suppose that the person in question is grossly ignorant of certain historical facts, historical facts which make it overwhelmingly likely that the Holocaust occurred. Indeed, perhaps the evidence that the Holocaust denier possesses is sufficiently impoverished and misleading (the misleading testimony provided by parents whom he had a default entitlement to trust; the propaganda to which he has been subjected, etc.) that his belief that the Holocaust never occurred is a perfectly reasonable thing for him to think, both objectively and by my own lights. His problem is not irrationality but ignorance. One might have thought that his gross ignorance is certainly something that I should take into account in evaluating the epistemic credentials of his belief. (Recall that, for reasons given above, evaluating a person's belief in the sense relevant to Independence* must go beyond merely making a judgment about the epistemic status of his belief given his evidence.) However, there seems to be a problem with my doing this. Suppose that it turns out that (as is plausible enough) the historical facts of which he is ignorant are the very same facts on which I base my own belief that the Holocaust occurred. In that case, in evaluating his belief, I should bracket my own assessment of these considerations. That is, I should set aside my own judgment that these considerations strongly support the view that the Holocaust occurred. But the problem then is this: my judgment that Holocaust denier is grossly ignorant when it comes to matters relating to the Holocaust is not at all independent of my assessment that the relevant considerations strongly confirm the occurrence of the Holocaust. That is, if I set aside my assessment that these facts strongly confirm the occurrence of the Holocaust, then I would no longer take someone's

ignorance of them to be a handicap in judging whether the Holocaust occurred. After all, there are ever so many facts ignorance of which I take to be no handicap at all when it comes to judging whether the Holocaust occurred. It is only because I judge that these facts confirm that the Holocaust occurred, that I take ignorance of them to be at all relevant to “the epistemic credentials” of someone’s belief about the Holocaust.

The worry then, is that once I engage in the relevant kind of bracketing, I will not have the resources needed to offer a suitably negative evaluation of the Holocaust denier’s belief. I think that there is a genuine problem here. But I don’t want to rest too much on the claim that this is a straightforward counterexample to principles like Independence(*). In the course of defending Independence against objections raised by Ernest Sosa (2010) and Jennifer Lackey (2010), Christensen (forthcoming) displays great ingenuity in parrying apparent counterexamples. Perhaps he or some other Conciliationist could show that there is some way in which I could reach a sufficiently negative assessment of the Holocaust denier’s belief even while bracketing my assessment of those considerations that underwrite my own, opposite belief. Even in that event, however, I think that we should be quite suspicious of the suggestion that the kind of bracketing exercise envisaged by Christensen plays an important role in how I should take the Holocaust denier’s opinion into account. Offhand, it seems akin to the suggestion that, in a case in which I discover that I hold two inconsistent beliefs, I should evaluate the credentials of one belief while bracketing my assessment that I have overwhelming evidence for the other.

The suspicion that Independence(*) is too strong might be enhanced when we recall the considerations which Christensen cites as motivation for adopting such principles in the first place:

The motivation behind the principle is obvious: it’s intended to prevent blatantly question-begging dismissals of the evidence provided by the disagreement of the others. It attempts to capture what would be wrong with a P-believer saying, e.g., “Well, so and so disagrees with me about p. But since P is true, she’s wrong about p. So however reliable she may generally be, I needn’t take her disagreement about p as any reason at all to change my belief.

There is clearly something worrisome about this sort of response to the disagreement of others. Used as a general tactic, it would seem to allow a non-expert to dismiss even the disagreement of large numbers of those he took to be experts in

the field.

An observation: the reasoning that Christensen describes in this passage is really quite awful. It is paradigmatically *dogmatic*, in the pejorative sense of that term. In view of the awfulness of such reasoning, it would be rather surprising, I think, if we needed to invoke a principle as strong as Independence or Independence* in order to explain what's wrong with it. That is, it would be surprising if we needed to adopt a principle that makes it obscure how I can legitimately appeal to my historical knowledge in evaluating the Holocaust denier's belief, in order to explain (e.g.) what would be wrong with my dogmatically dismissing a consensus of experts, or a person of arbitrarily great reliability, on the basis of my own, non-expert opinion.

This observation suggests a strategy that I will pursue in what follows. I'll attempt to show how someone who is either agnostic about strong Independence principles or who (like me) thinks that such principles are false can account for the badness of the relevant kind of reasoning in a perfectly natural way. This will leave such principles unmotivated, and, in view of their strength and the problems that they face, leave us with good reason to think that they are false.

Suppose that I arrive at the belief that p; I then hear you sincerely assert your belief that not-p. On the basis of your impressive track record, I know that you're generally reliable about this sort of question. Imagine the speech mentioned by Christensen in my mouth, addressed to you:

Well, you disagree with me about p. But since p is true, you must be wrong about p. So even though you're very reliable in general, I needn't take your disagreement about p as any reason at all to change my mind.

When I appeal to my belief to completely dismiss your contrary opinion, I am in effect inferring, from my belief that p, that your sincere assertion that not-p is *misleading evidence* about the truth of p. What account could someone who eschews appeal to strong Independence principles give of what is wrong with such reasoning?

Consider first a case in which my initial belief that p is unreasonable (that is, I unreasonably believe that p even before learning that you think otherwise). When I later infer that your testimony that not-p is misleading, my procedure is unreasonable and

dogmatic, inasmuch as I lack a reasonable basis for drawing the relevant conclusion: the proposition *p*, from which I infer the misleadingness of your testimony, is not among the things that I believe reasonably.¹¹

Of course, even if my belief was reasonable prior to receiving your testimony, this does not mean that I can then reasonably infer your testimony is misleading once I receive it. For even if I am initially justified in believing *p*, your testimony that not-*p* might undermine my justification, in which case I'm in no position to reasonably conclude that your testimony is misleading. Indeed, as Kripke (1971) emphasized, even if one initially *knows* that *p*, it might be unreasonable and dogmatic to dismiss subsequently encountered considerations that suggest that not-*p*. For once one is presented with those considerations, one might no longer know that *p* and thus no longer be in a position to rationally infer that the considerations are misleading (Harman 1973). Of course, if one once knew that *p*, then *p* is true, so the considerations that suggest that *p* is false must be misleading. But one is in no position to reasonably conclude this, once one's belief has been undermined.

So here is the short story about why it will often be unreasonable and dogmatic for me to dismiss your contrary opinion in the envisaged way: after I add the fact that you believe as you do to my stock of evidence, it will no longer be reasonable for me to believe that *p*, given what is then my total evidence. And if it's no longer reasonable for me to believe that *p*, then I lack any rational basis for inferring that your sincere testimony is misleading evidence. Of course, at one level Conciliationists will agree that that's the correct story; they will take themselves to have provided a deeper explanation. My current point is the modest one that the story outlined here about why the relevant kind of dogmatism is bad reasoning is certainly not the exclusive property of the Conciliationist; it can be told just as easily by the proponent of the Total Evidence View, and by others who either reject or are agnostic about Independence(*).

At this point, it's worth bearing in mind that there are cases in which it *is* reasonable for one to discount genuine evidence on the grounds that it is misleading (Sorensen 1988;

¹¹ In the usual case, if it's unreasonable for me to believe *p* prior to receiving your testimony that not-*p*, then it will still be unreasonable for me to believe *p* after receiving your testimony. There are unusual, trick cases in which this condition fails to hold, but I will ignore them here.

Kelly 2008). Consider the following extreme case:

TRUE STORY. I live with my family at 76 Alexander Street. On a fairly regular basis, we receive mail for a person named “Frederick Jacobs” at this address. This mail provides genuine evidence that someone named Jacobs lives at 76 Alexander Street. (Consider: when a passerby on the street, curious about who lives at this address, opens our mailbox and finds mail addressed to Jacobs, this increases the credibility of the relevant proposition for the passerby.) Nevertheless, on the basis of my knowledge that only members of my family live at 76 Alexander St. and that Jacobs is not a member of my family, I reasonably conclude that this evidence is misleading and dismiss it without further ado.¹²

Why isn’t my behavior in TRUE STORY unreasonable and dogmatic? Answer: given all of the evidence available to me which bears on the question of whether someone named Jacobs lives in my house—including those considerations that suggest that he does—it’s still reasonable for me to believe that he does not, and thus, to conclude that those considerations are misleading. This is what it’s reasonable for me to believe, given my total evidence.

Notice that (with a bit of stretching) we might even construe TRUE STORY as involving a piece of testimony: given the relevant conventions, an envelope that displays the name ‘Jacobs’ immediately above the address ‘76 Alexander St’ is a kind of written testimony which in this case constitutes misleading evidence for a false proposition. Even if we think of it in this way, however, it would certainly be foolish for me to bracket my assessment of all of the evidence which I take to underwrite my belief that no one named Jacobs lives in my house, in considering how to adjust my credence in the light of this piece of evidence.

The case of the Holocaust denier should be understood along these lines, I think. After I add the fact that he believes as he does to what I know, it’s still reasonable for me to have tremendous confidence that the Holocaust occurred, and (therefore) to infer that his belief is false. Contrast Christensen’s example in which I take my own non-expert opinion as a basis for concluding that all of the many experts who disagree with me are mistaken. The reason why this will typically be dogmatic and unreasonable is simply the following: even if my belief was initially reasonable, it will typically not be reasonable

¹² Although all of the details of the example are nonfictional, the inspiration for using them in this way is due to Crispin Wright (2004).

once I learn that all of the experts think otherwise, given how that will affect my epistemic situation. And given that it will no longer be reasonable for me to hold my original opinion, I will no longer have any legitimate basis from which to infer that the experts are mistaken.

So if the thought that is supposed to motivate Independence(*) is that we need some such principle in order to block dogmatism, or to account for why dogmatic reasoning is bad, I don't see it. Someone who rejects such principles can still account for the badness of intuitively dogmatic reasoning simply by appealing directly to the normative requirement that one take into account one's total evidence (as opposed to some proper subset of one's total evidence). In short, the Principle of Total Evidence can do all of the work that needs doing.

Objection: But consider cases such as Christensen's restaurant case, in which one's original evidence literally *entails* the correct answer. If one's original evidence entails that *p*, then it seems like one's total evidence will *always* support the belief that *p*, no matter how much misleading testimonial evidence one subsequently acquires, so long as that original evidence remains part of the total set. So it looks like we do need to invoke a principle like Independence after all, in order to allow such misleading counterevidence to change (at least eventually) what it is reasonable for one to believe about *p*.

Reply: there is a genuine puzzle here, but it is a mistake to think that that puzzle motivates the adoption of Independence or Independence*. After all, as Christensen himself would be the first to agree¹³, whenever one performs a non-trivial calculation, one should not be perfectly confident of one's answer even *before* another person comes on the scene (given one's awareness that one is fallible, etc.) But once it is granted that one should not be perfectly confident even before one's view is contradicted by a peer, there is no *additional* mystery or formal difficulty as to how acquiring that misleading testimonial evidence can push the credence that it is reasonable for one to have still lower. Of course, it is difficult to make sense of the idea that someone who possesses entailing evidence should invest less than maximal credence in the entailed proposition; indeed, orthodox theories of evidential probability would seem to rule this out (at least in cases in which the person is certain of the entailing evidence itself). But if, as both

¹³ See especially his 2007b.

Christensen and I think, this is a genuine phenomenon, then there is presumably some story to be told here. The crucial point is that there is no reason to think that the story in question entails Independence or Independence*, since those principles explicitly concern how one should assess the beliefs of other people, and the phenomenon arises even before other people come on the scene.

4. Two views about dogmatism

I have argued that we do not need to appeal to a principle such as Independence* in order to account for the badness of dogmatic reasoning. What might lead someone to think otherwise? Here is a speculation.¹⁴ I believe that many Conciliationists—and many others as well—tend to think about the intellectual vice of dogmatism in a particular way, a way that’s quite natural and intuitive, but ultimately mistaken. (Having said that, I should also say at once that I think that the issues here run deep and I won’t be able to pursue them very far in this paper. But I do want to at least put them on the table.) My hypothesis is that many people subscribe, even if only implicitly, to a thesis that I would express as follows:

Dogmatism is a formal vice.

What is it for something to be a “formal” vice? Consider first a paradigm from the moral sphere: the vice of hypocrisy. Someone who demands that others conform to putative moral standards that he himself transgresses is guilty of hypocrisy. This is a moral failing even if the standards that he fails to meet are not genuine moral standards at all. Because of this, one can be in a position to correctly charge another with hypocrisy without entering into potentially messy and difficult to resolve issues about what the correct moral standards are. Indeed, one might realize, from the first person perspective,

¹⁴ Because I have focused on the views of Christensen to this point, I want to explicitly disavow the suggestion that I am offering a speculative diagnosis as to why *he* accepts Independence. Having discussed these issues over a number of years with many Conciliationists and their fellow travelers, however, I am confident that implicit acceptance of the picture that I am about to describe is at least sometimes an operative factor. See also the discussion of Elga’s (2007) “[no] bootstrapping” argument below.

that a given action would be hypocritical and treat this as a reason not to perform it. In neither the first person nor the third person case does one have to make a substantive moral judgment about the correctness of the agent's moral standards, violation of which would constitute hypocrisy on his part. Thus, justified judgments of hypocrisy do not generally presuppose the truth of substantive claims about morality.

It is tempting to think of the intellectual vice of dogmatism in a parallel way: that at least in the usual case, whether someone has behaved dogmatically, or whether one would be dogmatic in responding to the disagreement of another person in a certain way, is something that can be determined without relying on a substantive and potentially precarious judgment about all-things-considered reasonableness or justifiedness. But I think that at least in the interesting cases, our sense that is so is an illusion: dogmatism is not a formal vice, inasmuch as justified attributions of it typically presuppose the truth of substantive claims about what it is reasonable to believe and infer given the totality of evidence available to the believer.

As a way of getting at the general issue, consider the status of so-called "Moorean" responses to revisionary philosophical theorizing. You present me with certain metaphysical principles that you endorse and the arguments for them; taken together, the principles entail that there are no tables and chairs, a consequence of which you are well aware and embrace as an exciting philosophical discovery. I consider your principles and supporting arguments. I then reason as follows: "Well, the principles seem quite plausible, and I don't have an intellectually satisfying diagnosis of where the arguments go wrong (if in fact they do). But of course, there are table and chairs. Therefore, the theory is false". Is this dogmatic on my part?

Someone who thinks of dogmatism as a formal voice will be strongly inclined to answer in the affirmative. But I think that that's the wrong answer. Rather, as in the cases presented in section 1 above, there simply isn't enough information provided in the fiction for us to know whether this particular bit of Moorean reasoning on my part is reasonable or not. Everything depends on whether, having considered your theory and the arguments for it, it's still reasonable for me to believe that there are tables and chairs, given what is now my total evidence. If it's not currently reasonable for me to believe this, then the way in which I dismiss your theory and the arguments is objectionable. But

on the other hand, if it's still reasonable for me to believe that there are tables and chairs, then this piece of Moorean reasoning is not dogmatic. Again, it all depends on what it is reasonable for me to believe given my total evidence, and that's not something that is specified in the fiction, or something that we're in a position to figure out given what is specified. There is, I think, no general objection to Moorean reasoning of this kind on the grounds of dogmatism, although of course, certain instances of such reasoning are objectionable in virtue of being dogmatic. And that's because dogmatism, unlike hypocrisy, is not a formal vice.¹⁵

Return to the case of disagreement, and suppose that one *did* think of dogmatism as a formal vice. Notice that in case, there is strong pressure to wheel in some principle like Independence(*) in order to explain what's wrong with the kind of paradigmatically bad reasoning described by Christensen, in which one uses one's non-expert opinion as a basis for completely dismissing an expert consensus, or the view of some hyper-reliable individual. Specifically: the fact that the reasoning is dogmatic is supposed to be something that holds (and therefore, is in principle recognizable) independently of the fact that the belief from which such reasoning proceeds is unreasonable on the dogmatist's total evidence. So we need some principle that counts the reasoning as bad/dogmatic, *other than* the Requirement of Total Evidence. But once we give up on the idea that dogmatism is a formal vice, we give up on the project of identifying instances of dogmatism in a way that does not presuppose substantive judgments about what is and isn't reasonable to believe given the subject's total evidence.

Objection: "You yourself have assumed that Christensen's putative examples of dogmatic reasoning are genuine instances. But of course, Christensen never said anything about the subject's total evidence! So you too think that we can typically identify instances of dogmatism, independently of knowing what it's reasonable to believe given the subject's total evidence. That is, you yourself think that dogmatism is a formal vice after all".

¹⁵ In fact, I think that one effect of the tendency to think of dogmatism as a formal vice is that the strength of Moorean responses to revisionary philosophical theorizing is often greatly underestimated. For a defense of such reasoning, see Kelly (2005), Kelly (2008). For more general reflections on the possibility that dogmatism is not a formal vice and related matters, see my (2011).

Reply: It's true that Christensen never tells us about the subject's total evidence, or about all of the subject's total evidence that bears on *p*. But he does tell us that part of the subject's total evidence is made up of extremely substantial and strong evidence that not-*p* (the consensus of the experts that not-*p*, etc.). Given that this evidence is included in the subject's total evidence, one would (at the very least) have to tell an Extremely Unusual Story about the rest of the subject's total evidence, in order for it be reasonable for the subject to believe *p* in a way that would render the envisaged reasoning non-dogmatic. So naturally, when we think about the case, we implicitly fill in the details in a way that makes the reasoning dogmatic: that is, we implicitly assume that nothing like the Extremely Unusual Story is true in the fiction. (Compare the way in which, in judging that the protagonists in Gettier's paper lack knowledge, we assume such things as that their beliefs are not over-determined in ways that render them knowledge after all.)

Again, someone who denies that dogmatism is a formal vice will tend to think that there are *fewer* genuine epistemic norms than someone who assumes that it is a formal vice. In order to see how this plays out in a concrete case, consider what is perhaps the single most prominent argument for Conciliationism in the literature, Adam Elga's "no bootstrapping" argument for his version of the Equal Weight View. It is worth quoting in full:

Suppose that you and your friend are to judge the truth of a claim, based on the same batch of evidence. Initially, you count your friend as an epistemic peer—you think that she is about as good as you at judging the claim. In other words, you think that, conditional on a disagreement arising, the two of you are equally likely to be mistaken. Then the two of you perform your evaluations. As it happens, you become confident that the claim is true, and your friend becomes equally confident that it is false.

When you learn of your friend's opposing judgment, you should think that the two of you are equally likely to be correct. The reason is [this]. If it were reasonable for you to give your own evaluation extra weight—if it were reasonable to be more than 50% confident that you are right—then you would have gotten some evidence that you are a better evaluator than your friend. But that is absurd.

[T]he absurdity is made more apparent if we imagine that you and your friend evaluate the same long series of claims. Suppose for *reductio* that whenever the two of you disagree, you should be, say, 70% confident that your friend is the mistaken one. It follows that over the course of many disagreements, you should end up extremely confident that you have a better track record than your friend. As a result, you should end up extremely confident that you are a better evaluator. But that is absurd. Without some antecedent reason to think that you are a better evaluator, the

disagreements between you and your friend are no evidence that she has made most of the mistakes (2007: 487).

Elsewhere I have criticized this argument at some length (2007 sect.x); here I will concentrate on those aspects that intersect most directly with the present issue.

Elga takes the argument of this passage to successfully undermine any alternative to The Equal Weight View. In particular, he takes the argument offered here to undermine both what he calls ‘The Extra Weight View’—according to which each party to the dispute is permitted to give some special, presumptive weight to his or her own judgment—as well as views akin to The Total Evidence View, on which it potentially matters which of the parties has in fact done a better job evaluating the evidence.¹⁶ However, I believe that while the argument has considerable force against the former sort of view, it has little to none against the latter.

In order to see this, let’s focus our attention directly on the situation in which Elga claims the absurdity of any alternative to The Equal Weight View is most apparent, viz. the situation in which you and your friend each evaluate a long series of claims. Elga formulates the argument as a *reductio ad absurdum*. The supposition from which the absurd consequences are alleged to follow is this:

whenever you and your friend disagree, you should be, say, 70% confident that your friend is the mistaken one

The crucial fact here is the following: this supposition is *not* something to which a proponent of The Total Evidence View is committed. That is, the proponent of The Total Evidence View is not committed to the idea that, whenever you and your friend disagree, you should be n% confident that your friend is the one who has made the mistake (where n is some number greater than 50). Indeed, on the contrary: the proponent of The Total Evidence View will stand with Elga in rejecting any such general policy as an unreasonable one. On The Total Evidence View, it’s not true, in general, that you should

¹⁶ Elga makes the last point explicit on the same page:

Again, this absurdity is independent of who has in fact evaluated the claims properly. Even if in fact you have done a much better job than your friend at evaluating the claims, simply comparing your verdicts to those of your friend gives you no evidence that this is so (p.487).

be more confident that your friend has made the mistake whenever the two of you disagree. Nor is there some general answer to the question of how confident you should be that it's your friend who has made the mistake (as there is on both The Extra Weight View and on The Equal Weight View). And this is because how confident it's reasonable for you to be that your friend has made a mistake is not something that floats entirely free of the evidence on which he bases his opinion. Thus, since the proponent of The Total Evidence View would not accept the supposition from which Elga derives the absurd consequence, the *reductio ad absurdum* on offer cannot show that her view is false.

In giving this argument, Elga simply assumes that there must be a general norm which has something like the following form:

In any case of peer disagreement, you should conclude that your peer's opinion is $n\%$ likely to be correct, and revise your own opinion accordingly.

Notice that a norm of this form will say nothing at all about the reasonableness or justifiedness of anyone's opinion, or about what the total evidence supports (either before or after the disagreement is discovered). The assumption that there is a genuine norm of this kind is thus a very substantial one. Of course, once one makes the assumption that there is such a norm, the suggestion that the correct value for n is '50' is absolutely irresistible. Certainly, given what it is to count someone as an epistemic peer, it would be completely bizarre to suspect that some number other than 50 might be the real value of n . But there is another possibility: namely, that how one is rationally required to respond to a disagreement is not typically something that is fixed independently of substantive normative facts about how well-supported one's original view is. On the alternative picture, how confident one is rationally permitted to be that some proposition is true upon discovering that a peer thinks otherwise might vary significantly from case to case. To take this alternative picture seriously is to take seriously the possibility that there is no genuine norm of belief revision that has the same general form as the one endorsed by Elga.

Of course, in some ways it would make our cognitive lives much easier if there were genuine norms of the relevant kind. Upon learning that a peer disagrees with an opinion that one holds, one would learn that one is now rationally required to invest a certain

credence in the relevant proposition. Simply on the basis of applying the norm to the original credences, one would be in a position to know what credence one is rationally required to have now. That is, one would be in a position to know what credence one is rationally required to have, *without* needing to make a substantive judgment about what it is reasonable for one to believe given one's total evidence, a type of judgment that is highly fallible, especially in the kind of "hard cases" that are salient in this context.¹⁷ So it would, I think, be much easier to figure out what one is rationally required to believe if there really were norms of the relevant kind. But the fact that it would make our cognitive lives easier if there were such norms is not itself a good reason to think that they exist. Faced with a peer who disagrees, knowing how one is rationally required to respond will typically require an extremely substantive judgment about one's overall epistemic situation, as opposed to the straightforward application of a general norm that dictates agnosticism in all such cases. Such are the burdens of judgment.¹⁸

¹⁷ To be sure, someone who has an Elga-like picture will readily agree that the fact that such-and-such an attitude (in the canonical case, agnosticism) is rationally required is a matter of it's being rationally required *given one's total evidence*. (I'm certainly not accusing such a person of having to deny the Requirement of Total Evidence.) The point is rather that the judgment that *agnosticism is rationally required given one's total evidence* is in effect trivial and not substantive: for the fact that one's total evidence now includes a particular piece of information about credences suffices, in connection with the norm, to determine that agnosticism is the rationally required attitude.

¹⁸ Earlier versions of this paper were presented at New York University, at the Pontificia Universidade Catolica de Rio Grande do Sul in Porto Alegre, Brazil, and at a meeting of Fritz Warfield's graduate seminar at the University of Notre Dame; I am grateful to those audiences for their feedback. Thanks also to David Christensen and Nate King for written comments on earlier drafts .

References

- Christensen, David (2007a). "Epistemology of Disagreement: the Good News", in *The Philosophical Review*, volume 116, no.2, pp.187-217.
- Christensen, David (2007b). "Does Murphy's Law Apply in Epistemology? (2007b) Self-Doubt and Rational Ideals," *Oxford Studies in Epistemology* 2 (2007): 3 - 31.
- Christensen, David (2009). "Disagreement as Evidence: The Epistemology of Controversy". *Philosophy Compass* 4/5: 756-767.
- Christensen, David (2010). "Higher Order Evidence", in *Philosophy and Phenomenological Research*. 81 (1): 185-215
- Christensen, David (forthcoming) "Disagreement, Question-Begging and Epistemic Self-Criticism". In *Philosophers Imprint*.
- Cohen, Stewart (this volume). "A Defense of the (Approximately) Equal Weight View".
- Elga, Adam (2007). "Reflection and Disagreement", in *Nous*, vol.XLI (No.3): 478-502.
- Elga, Adam (2010). "How to Disagree About How to Disagree". In Feldman and Warfield (eds.) *Disagreement* (Oxford University Press, 2010).
- Feldman, Richard (2003). *Epistemology*. (Prentice Hall, Upper Saddle River, NJ).
- Feldman, Richard (2006). "Epistemological Puzzles About Disagreement" in Stephen Hetherington (ed.) *Epistemology Futures* (Oxford University Press, Oxford), pp.216-236.
- Feldman, Richard (2007). "Reasonable Religious Disagreements", in Louise Antony (ed.) *Philosophers Without Gods: Meditations on Atheism and the Secular Life* (Oxford: Oxford University Press), pp.194-214.
- Harman, Gilbert (1973). *Thought*. (Princeton, NJ: Princeton University Press).
- Jehle, David and Fitelson, Branden (2009). "What is the 'Equal Weight View'?" *Episteme* 6: 280 – 293.
- Kelly, Thomas (2005). "Moorean Facts and Belief Revision: Can the Skeptic Win?" In John Hawthorne (ed.) *Philosophical Perspectives, vol.19: Epistemology* (Blackwell): 179-209.
- Kelly, Thomas (2007). "Peer Disagreement and Higher Order Evidence". Official version online at <http://www.princeton.edu/~tkelly/papers>. Also in Feldman and Warfield (eds.) *Disagreement* (Oxford University Press, 2010); as well as in Goldman and Whitcomb

(eds.) *Social Epistemology* (Oxford University Press).

Kelly, Thomas (2008). "Common Sense as Evidence: Against Revisionary Ontology and Skepticism", in Peter French (ed.) *Midwest Studies in Philosophy*, vol. XXXII: 'Truth and Its Deformities' (Blackwell).

Kelly, Thomas (2011). "Following the Argument Where It Leads". *Philosophical Studies*.

Kelly, Thomas (manuscript). "Believers as Thermometers".

Kornblith, Hilary (2010). "Belief in the Face of Controversy". In Feldman and Warfield (eds.) *Disagreement* (Oxford University Press).

Kripke, Saul (1971). "On Two Paradoxes of Knowledge". Lecture delivered to the Cambridge Moral Sciences Club. Transcribed by G.E.M. Anscombe.

Lackey, Jennifer (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.

Pryor, James (2000). "The Skeptic and the Dogmatist" *Noûs* 34 (4):517–549.

Sorensen, Roy (1988). "Dogmatism, Junk Knowledge, and Conditionals", *The Philosophical Quarterly* XXXVIII 153 (1988): 433-454.

Sosa, Ernest (2010). "The Epistemology of Disagreement", in A. Haddock, A. Millar and D. Pritchard, eds. *Social Epistemology* (Oxford: Oxford University Press): 278-297.

Wright, Crispin (2004). "Wittgensteinian Certainties". In *Wittgenstein and Scepticism*, ed. Denis McManus (Oxford: Routledge): 22-54.