

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

Peer Disagreement and Higher Order Evidence¹

Thomas Kelly
Princeton University

1. Introduction

My aim in this paper is to develop and defend a novel answer to a question that has recently generated a considerable amount of controversy. The question concerns the normative significance of *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 which is available to you but not to me, or vice versa. For the sake of concreteness, we might picture

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.

You and I are weather forecasters attempting to determine whether it will rain tomorrow. We both have access to the same meteorological data.

You and I are professional philosophers interested in the question of whether free will is compatible with determinism. Each of us is thoroughly acquainted with all of the extant arguments, thought experiments, and intuition pumps that the literature has to offer.

¹This paper is something of a sequel to Kelly (2005). While in many respects it is faithful to the position advanced there, it departs in others; significant departures are noted along the way. Earlier versions of this paper were presented at NYU, MIT, Rutgers University and at the University of California at Irvine; I am grateful to the audiences present on those occasions. In addition, I would like to thank Hartry Field, Jim Pryor, Ernest Sosa, Roy Sorensen, Walter Sinnott-Armstrong, Jim Joyce, Philip Pettit, Sarah McGrath, Adam Elga, David Christensen, Margaret Gilbert, Aaron James, Allan Gibbard, Daniel Greco and Aaron Bronfman for helpful conversations on the topic.

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 us suppose that, to the extent that we do possess evidence about who is more reliable—evidence afforded, perhaps, by a comparison of our past track records—such evidence suggests that we are more or less equally reliable when it comes to making judgments about the domain in question.² Nevertheless, despite being 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, or that it will rain tomorrow, or that free will and determinism are compatible, while you find yourself equally confident of the opposite. Question: once you and I learn that the other has arrived at a different conclusion despite having been exposed to the same evidence and arguments, how (if at all) should we revise our original views?

Some philosophers hold that in such circumstances, you and I are rationally required to *split the difference*. According to this line of thought, it would be unreasonable for either of us to simply retain his or her original opinion. Indeed, given the relevant symmetries, each of us should give *equal* weight to his or her opinion and to the opinion of the other in arriving at a revised view. Thus, given that I am confident that the accused is guilty while you are equally confident that he is not, both of us should retreat to a state of agnosticism in which we suspend judgment about the question. This is

The Equal Weight View: In cases of peer disagreement, one should give equal weight to the opinion of a peer and to one's own opinion.

Recently, The Equal Weight View has been endorsed by a number of philosophers. Here, for example, is Richard Feldman:

² Of course, the kind of uncontroversial 'track record' evidence that bears most directly on questions of comparative reliability will be much easier to come by in some domains than in others. (In this respect, contrast reliability in accurately forecasting the weather and reliability in accurately answering metaphysical questions.)

[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 judgement is called for (2006, p.235).³

It is no surprise that The Equal Weight View has found sophisticated advocates; it is in many respects an appealing view. Indeed, reflection on certain kinds of cases can make it seem almost trivial or obviously true. Consider, for example, cases involving conflicting perceptual judgments such as the following:

Case 1. You and I, two equally attentive and well-sighted individuals, stand side-by-side at the finish line of a horse race. The race is extremely close. At time t_0 , just as the first horses cross the finish line, it looks to me as though Horse A has won the race in virtue of finishing slightly ahead of Horse B; on the other hand, it looks to you as though Horse B has won in virtue of finishing slightly ahead of Horse A. At time t_1 , an instant later, we discover that we disagree about which horse has won the race. How, if at all, should we revise our original judgments on the basis of this new information?

Many find it obvious that, in such circumstances, I should abandon my original view that Horse A won the race and you should abandon your original view that Horse B won the race. We should become agnostics about which horse won the race until further evidence becomes available. This, of course, is exactly what The Equal Weight View enjoins. But one might expect that what holds for perceptual judgments holds also for judgments of other kinds, and thus, in general.

³Compare Feldman (2003), which, after reviewing a number of examples of the kind at issue here, draws the conclusion that 'In the situations most plausibly thought to be cases of reasonable disagreement, suspension of judgment is the reasonable attitude to take toward the disputed proposition' (p.189).

The Equal Weight View is explicitly embraced by Adam Elga (forthcoming), whose views I consider at some length below; David Christensen (2007) exhibits considerable sympathy for a policy of 'splitting the difference' throughout his own discussion of the topic. Although the view that I will put forth differs from theirs, I have learned much from each of these authors.

Further evidence for The Equal Weight View seems to be afforded by certain natural analogies involving inanimate measuring devices. Consider for example

Case 2. You and I are each attempting to determine the current temperature by consulting our own personal thermometers. In the past, the two thermometers have been equally reliable. At time t_0 , I consult my thermometer, find that it reads '68 degrees', and so immediately take up the corresponding belief. Meanwhile, you consult your thermometer, find that it reads '72 degrees', and so immediately take up that belief. At time t_1 , you and I compare notes and discover that our thermometers have disagreed. How, if at all, should we revise our original opinions about the temperature in the light of this new information?⁴

I take it as obvious that in these circumstances I should abandon my belief that it is 68 degrees and you should abandon your belief that it is 72 degrees. In particular, it would be unreasonable for me to retain my original belief simply because this was what *my* thermometer indicated. Indeed, inasmuch as the relevant evidence available to us is exhausted by the readings of the two thermometers, neither of us should be any more confident of what his thermometer says than of what the other person's thermometer says. In these circumstances, we should treat the conflicting thermometer readings as equally strong pieces of evidence. But--one might naturally conclude--what holds for the conflicting readings of equally reliable thermometers holds also for the conflicting judgments of individuals who are peers in the relevant respects. The mere fact that I originally judged that the accused is guilty is no reason for me to retain that view once I learn that you originally judged that he is innocent. Just as I should retreat to a state of agnosticism about whether the temperature is 68 or 72 degrees once I learn what your thermometer indicates, so too I should retreat to a state of agnosticism about whether the accused is guilty or innocent once I learn your opinion about the matter.

In view of considerations such as these and others that have been offered on its behalf, The Equal Weight View can seem quite compelling. Nevertheless, I believe that here appearances are misleading: The Equal Weight View is false. The main negative burden of what follows is to show that (and why) this is so. After offering a critique of The

⁴ A case of this general form was put to me by Roy Sorensen in conversation. Compare Christensen's (2007, p.196) 'Acme watch' example and Feldman (2006, p.234).

Equal Weight View, I will use that critique as a point of departure for the development of an alternative proposal about how we should respond to peer disagreement. For reasons that will emerge, I call this alternative proposal **The Total Evidence View**.

I begin with some taxonomy.

Philosophers who hold views inconsistent with The Equal Weight View maintain that, in at least some cases of peer disagreement, it can be reasonable to *stick to one's guns*.⁵ A particularly radical alternative is this:

The No Independent Weight View: In at least some cases of peer disagreement, it can be perfectly reasonable to give no weight at all to the opinion of the other party.

That is, even if one retains one's original opinion with wholly undiminished confidence upon learning that a peer thinks otherwise, one's doing so might be perfectly reasonable.

According to more moderate alternatives, while one is always rationally required to give at least some weight to the opinion of a peer, one is not always required to split the difference. That is, even if one's new opinion is closer to one's original opinion than to the original opinion of one's peer, one's new opinion might nevertheless be perfectly reasonable. Of course, there are many possible views of this kind. We might picture these possibilities as constituting a spectrum: at one end of the spectrum sits The Equal Weight View; at the other end, The No Independent Weight View; in between, the more moderate alternatives, arranged by how much weight they would have one give to the opinion of a peer relative to one's own. The more weight one is required to give to a peer's opinion relative to one's own, the more the view in question will resemble The Equal Weight View; the less weight one is required to give, the more it will resemble The No Independent Weight View.

Among alternatives to The Equal Weight View, another distinction is worth marking. Suppose that, upon learning that we hold different opinions about some issue, neither you nor I splits the difference: each of us either simply retains his or her original opinion, or else moves to a new opinion that is closer to that opinion than to the original opinion of the other. Again, according to The Equal Weight View, both you and I are unreasonable

⁵Notable here are van Inwagen (1996), Plantinga (2000a, 2000b), and Rosen (2001); another is Kelly (2005).

for responding to our disagreement in this way. Among views inconsistent with The Equal Weight View, distinguish between those according to which you and I might *both* be reasonable in responding in this way and those according to which *at most one* of us is being reasonable. As an example of the former, consider a view according to which everyone is rationally entitled to give some special, presumptive weight to his or her own judgment.⁶ If such a view is true, then both you and I might be perfectly reasonable even though neither one of us splits the difference. As an example of the latter kind of view, consider a view according to which how far you and I should move in response to our disagreement depends on whose original opinion better reflects our original evidence (Kelly 2005). Given such a view, and given certain further assumptions, it might be that when you and I fail to split the difference, at most one of us is being reasonable.

Taking these two distinctions together, the view most radically at odds with The Equal Weight View would seem to be the following:

The Symmetrical No Independent Weight View: In at least some cases of peer disagreement, both parties to the dispute might be perfectly reasonable even if neither gives any weight at all to the opinion of the other party.

Thus, according to The Symmetrical No Independent Weight View, even if both you and I remain utterly unmoved upon learning that the other holds a different opinion, it might be that neither one of us is responding unreasonably.

It is not my purpose to defend The Symmetrical No Independent Weight View. Indeed, the view about peer disagreement that I will ultimately endorse is consistent with both it and its negation. That having been said, I am inclined to think that The Symmetrical No Independent Weight View is true. Moreover, I also believe that, precisely because it contrasts so sharply with The Equal Weight View, considering it can help to illuminate the latter by making plain some of the less obvious dialectical commitments incurred by proponents of The Equal Weight View. For these reasons, I want to briefly explore what might be said on its behalf.

⁶ Compare ‘The Extra Weight View’ discussed by Elga (forthcoming) who argues against it.

2. Cases in Which Both You and I are Perfectly Reasonable, Despite Giving No Weight to the Other's Point of View

First, a preliminary remark about The Equal Weight View. The Equal Weight View is sometimes defended in contexts in which the propositional attitude of belief is treated as an all-or-nothing matter: for any proposition that one considers, one has in effect three doxastic options--one either believes the proposition, disbelieves the proposition, or suspends judgment as to its truth.⁷ However, in considering The Equal Weight View, it is for various reasons more natural to treat belief not as an all-or-nothing matter but rather as a matter of degree. Indeed, it does not seem that The Equal Weight View can even be applied in full generality in a framework which treats belief as an all-or-nothing matter. Thus, consider a possible world which consists of two peers, one of whom is a theist and the other of whom is an atheist. When the theist and the atheist encounter one another, the response mandated by The Equal Weight View is clear enough: the two should split the difference and become agnostics with respect to the question of whether God exists. Suppose, however, that the two person world consists not of a theist and an atheist but rather an atheist and an agnostic. How do they split the difference? (In this case, of course, agnosticism hardly represents a suitable compromise.) In general, the simple tripartite division between belief, disbelief and suspension of judgment does not have enough structure to capture the import of The Equal Weight View when the relevant difference in opinion is that between belief and suspension of judgment, or between suspension of judgment and disbelief. Clearly, the natural move at this point is to employ a framework which recognizes more fine-grained psychological states. Let us then adopt the standard Bayesian convention according to which the credence which one invests in a given proposition is assigned a numerical value between 0 and 1 inclusive, where 1 represents maximal confidence that the proposition is true, 0 represents maximal confidence that the proposition is false, .5 represents a state of perfect agnosticism as to the truth of the proposition, and so on. Thus, if the agnostic gives credence .5 to the proposition that God exists while the atheist gives credence .1 to the same proposition,

⁷See, for example, Feldman (2003, 2006).

the import of The Equal Weight View is clear: upon learning of the other's opinion, each should give credence .3 to the proposition that God exists.

Moreover, even if one restricts one's attention to what are sometimes called 'strong disagreements', i.e., cases in which the relevant proposition is initially either believed or disbelieved by the parties,⁸ it seems that an advocate of The Equal Weight View still has strong reasons to insist on a framework which treats belief as a matter of degree. For consider a world of three peers, two of whom are theists and one of whom is an atheist. The animating thought behind The Equal Weight View, viz. that the opinion of any peer should count for no more and no less than that of any other, would seem to be clearly violated by the suggestion that the parties to the dispute should retreat to a state of agnosticism, since that would seem to give more weight to the opinion of the atheist than to the opinion of either theist. (The atheist's opinion is in effect given as much weight as the opinions of both theists taken together in determining what should ultimately be believed by the three.) On the other hand, the suggestion that theism wins simply because the atheist finds himself outnumbered would seem to give too little weight to the atheist's original opinion if it is understood to mean that all three should ultimately end up where the two theists begin. Once again, it seems that an advocate of The Equal Weight View should insist on a framework which treats belief as a matter of degree since only such a framework can adequately capture what is clearly in the spirit of his or her view.

Having noted this elementary point, I will now describe a possible case in which you and I are both perfectly reasonable despite giving zero weight to the other person's opinion:

Case 3. How things stand with me:

At time t_0 , my total evidence with respect to some hypothesis H consists of E . My credence for H stands at .7. Given evidence E , this credence is perfectly reasonable. Moreover, if I was slightly less confident that H is true, I would also be perfectly reasonable. Indeed, I recognize that this is so: if I met someone who shared my evidence but was slightly less confident that H was true, I would not consider that person unreasonable for believing as she does.

⁸ Again, this is characteristic of Feldman's work on the topic.

How things stand with you:

At time t_0 , your total evidence with respect to H is also E. Your credence for H is slightly lower than .7. Given evidence E, this credence is perfectly reasonable. Moreover, you recognize that, if your credence was slightly higher (say, .7), you would still be perfectly reasonable. If you met someone who shared your evidence but was slightly more confident that H was true, you would not consider that person unreasonable for believing as she does.

At time t_1 , we meet and compare notes. How, if at all, should we revise our opinions?

According to The Equal Weight View, you are rationally required to increase your credence while I am rationally required to decrease mine. But that seems wrong. After all, *ex hypothesi*, the opinion that I hold about H is within the range of perfectly reasonable opinion, as is the opinion that you hold. Moreover, both of us have recognized this all along. Why then would we be rationally required to change?

One sympathetic to The Equal Weight View might attempt to heroically defend the idea that you and I are rationally required to revise our original credences in these circumstances. However, a more promising line of resistance, I think, is to deny that Case 3 is possible at all. That is, an adherent of The Equal Weight View should endorse

The Uniqueness Thesis: For a given body of evidence and a given proposition, there is some one level of confidence that it is uniquely rational to have in that proposition given that evidence.⁹

Suppose that The Uniqueness Thesis is true. Then, if it is in fact reasonable for me to give credence .7 to the hypothesis, it follows that you are guilty of unreasonable diffidence for being even slightly less confident. On the other hand, if you are reasonable in being slightly less confident than I am, then I am guilty of being unreasonably

⁹ 'The Uniqueness Thesis' is Feldman's (2007) label; compare Christensen's (2007) 'Rational Uniqueness'. Feldman both argues for and endorses the thesis; Christensen exhibits some sympathy for the thesis and offers some considerations for thinking that it is true. White (2005) argues for the thesis at length but stops short of endorsing it.

overconfident. Hence, the description of Case 3 offered above is incoherent; Case 3 is not in fact a possible case.

Clearly, The Uniqueness Thesis is an extremely strong claim: for any given batch of evidence, there is some one correct way of responding to that evidence, any slight departure from which already constitutes a departure from perfect rationality. How plausible is The Uniqueness Thesis? For my part, I find that its intuitive plausibility depends a great deal on how we think of the psychological states to which it is taken to apply. The Uniqueness Thesis seems most plausible when we think of belief in a maximally coarse-grained way, as an all-or-nothing matter.¹⁰ On the other hand, as we think of belief in an increasingly fine-grained way, the more counterintuitive it seems. But as we have seen, the advocate of The Equal Weight View has strong reasons to insist on a framework which employs a fine-grained notion of belief.

Some philosophers find it pre-theoretically obvious that The Uniqueness Thesis is false.¹¹ Many others accept substantive epistemological views from which its falsity follows.¹² Although The Uniqueness Thesis is inconsistent with many popular views in epistemology and the philosophy of science, its extreme character is perhaps best appreciated in a Bayesian framework. In Bayesian terms, The Uniqueness Thesis is equivalent to the suggestion that there is some single prior probability distribution that it is rational for one to have, any slight deviation from which already constitutes a departure

¹⁰Most plausible, but still not especially plausible, I think. Again, it comes under pressure from marginal cases. Suppose that the evidence available to me is just barely sufficient to justify my belief that it will rain tomorrow: if the evidence was even slightly weaker than it is, then I would be unjustified in thinking that it will rain. Suppose further that you have the same evidence but are slightly more cautious than I am, and so do not yet believe that it will rain tomorrow. It is not that you are dogmatically averse to concluding that it will rain; indeed, we can suppose that if the evidence for rain gets even slightly stronger, then you too will take up the relevant belief. Is there some guarantee, given what has been said so far, that you are being less reasonable than I am?—I doubt it.

¹¹Here, for example, is Gideon Rosen:

It should be obvious that reasonable people can disagree, even when confronted with a single body of evidence. When a jury or a court is divided in a difficult case, the mere fact of disagreement does not mean that someone is being unreasonable (2001, p.71).

¹² See, e.g., the brief survey in White (2005) pp.445-446.

from perfect rationality. This contrasts most strongly with so-called ‘orthodox’ Bayesianism, according to which any prior probability distribution is reasonable so long as it is probabilistically coherent. Of course, many Bayesians think that orthodoxy is in this respect overly permissive. But notably, even those Bayesians who are considered Hard Liners for holding that there are substantive constraints on rational prior probability distributions other than mere probabilistic coherence typically want nothing to do with the suggestion there is some uniquely rational distribution. With respect to this long-running debate then, commitment to The Uniqueness Thesis yields a view that would be considered by many to be beyond the pale, too Hard Line even for the taste of the Hard Liners themselves.

Of course, despite its radical character, The Uniqueness Thesis might nevertheless be true. In fact, some formidable arguments have been offered on its behalf.¹³ Because I believe that The Uniqueness Thesis is false, I believe that The Symmetrical No Independent Weight View is true, and (therefore) that The Equal Weight View is false. However, especially in light of the fact that here I will neither address the arguments for The Uniqueness Thesis nor argue against it more directly, I will not appeal to the possibility of so-called ‘reasonable disagreements’ in arguing against The Equal Weight View. Indeed, because I am convinced that we should reject The Equal Weight View in any case, I will proceed in what follows as though (what I take to be) the fiction of uniqueness is true. My dialectical purpose in emphasizing the apparent link between The Uniqueness Thesis and The Equal Weight View is a relatively modest one. As noted above, The Equal Weight View can sometimes seem to be almost *obviously* or *trivially* true, as though its truth can be established by quick and easy generalization from a few simple examples or analogies. However, if I am correct in thinking that commitment to The Equal Weight View carries with it a commitment to The Uniqueness Thesis, then this is one possibility that can be safely ruled out. Even if turns out to be true, The Uniqueness Thesis is an extremely strong and unobvious claim. Inasmuch as the ultimate

¹³I take the most formidable case to have been made by White (2005), although he himself does not endorse the thesis. I respond to some, though not all, of White’s arguments in my “‘Epistemic Permissiveness’: Comments on White”, available at <supply url>.

tenability of The Equal Weight View is bound up with its ultimate tenability, The Equal Weight View is similarly an extremely strong and unobvious claim.

I turn next to some arguments against The Equal Weight View.

3. Why We Should Reject the Equal Weight View

Let us suppose for the sake of argument then, that The Uniqueness Thesis is correct: for a given batch of evidence, there is some one way of responding to that evidence that is the maximally rational way. Consider

Case 4. Despite having access to the same substantial body of evidence E, you and I arrive at very different opinions about some hypothesis H: while I am quite confident that H is true, you are quite confident that it is false. Indeed, at time t_0 , immediately before encountering one another, my credence for H stands at .8 while your credence stands at .2. At time t_1 , you and I meet and compare notes. How, if at all, should we revise our respective opinions?

According to The Equal Weight View, you and I should split the difference between our original opinions and each give credence .5 to H. As a general prescription, this strikes me as wrongheaded, for the following reason. Notice that, in the case as it has been described thus far, nothing whatsoever has been said about the relationship between E and H, and in particular, about the extent to which E supports or fails to support H. But it is implausible that how confident you and I should be that H is true at time t_1 is wholly independent of this fact. For example, here is a way of filling in the details of the case which makes it implausible to suppose that you are rationally required to split the difference with me:

Case 4, continued. In fact, hypothesis H is quite unlikely on evidence E. Your giving credence .2 to H is the reasonable response to that evidence. Moreover, you respond in this way precisely because you recognize that H is quite unlikely on E. On the other hand, my giving credence .8 to H is an unreasonable response and reflects the fact that I have significantly overestimated the probative force of E with respect to H.

At time t_0 then, prior to encountering the other person, things stand as follows: you hold a reasonable opinion about H on the basis of your total evidence while I hold an unreasonable opinion about H on the basis of the same total evidence. (Again, the difference in the normative statuses of our respective opinions is due to the fact that your opinion is justified by our common evidence while mine is not.) If one were to ask which one of us should revise his or her view at *this* point, the answer is clear and uncontroversial: while it is reasonable for you to retain your current level of confidence, I should significantly reduce mine, since, *ex hypothesi*, this is what a correct appreciation of my evidence would lead me to do.

For an advocate of The Equal Weight View, this seemingly important asymmetry completely washes out once we become aware of our disagreement. Each of us should split the difference between his or her original view (regardless of whether that view was reasonable or unreasonable) and the original view of the other (regardless of its status). I take this to be an extremely dubious consequence of The Equal Weight View. We should be clear, however, about exactly which consequences of The Equal Weight View warrant suspicion and which do not. According to The Equal Weight View, after you and I meet, I should be significantly less confident that the hypothesis is true. That much is surely correct. (After all, I should have been significantly less confident even before we met.) The Equal Weight View also implies that, after we meet, you should be more confident that the hypothesis is true, despite having responded correctly to our original evidence. While less obvious, this is also—for reasons that I explore below—not implausible. What *is* quite implausible, I think, is the suggestion that you and I are rationally required to make *equally* extensive revisions in our original opinions, given that your original opinion was, while mine was not, a reasonable response to our original evidence. After all, what it is reasonable for us to believe after we meet at time t_1 presumably depends upon the total evidence that we possess at that point. Let's call the total evidence that we possess at time t_1 E^* . What does E^* include? Presumably, E^* includes the following:

Our original body of evidence E

The fact that I responded to E by believing H to degree .8

The fact that you responded to E by believing H to degree .2.

Notice that, on The Equal Weight View, the bearing of E on H turns out to be completely irrelevant to the bearing of E* on H. In effect, what it is reasonable for you and I to believe about H at time t1 supervenes on how you and I respond to E at time t0. With respect to playing a role in determining what is reasonable for us to believe at time t1, E gets completely swamped by purely psychological facts about what you and I believe. (This despite the fact that, on any plausible view, it was highly relevant to determining what it was reasonable for us to believe back at time t0.) But why should the normative significance of E completely vanish in this way?

We can, of course, imagine a case in which it would be reasonable for one to form an opinion about H by simply splitting the difference between your opinion and mine: namely, a case in which those opinions are the only relevant evidence that one possesses. Imagine, for example, the position of a third party who lacks any direct access to E, and knows only that, of two equally well-informed parties, one gives credence .2 and the other gives credence .8 to hypothesis H. (Suppose also that the individual lacks any other relevant evidence.) For an individual so situated, assigning a probability of .5 to H is at least as reasonable as any other course. Perhaps the same would be true of you and I, if, at some still later time t2, we completely lost access to our original evidence--say, in virtue of forgetting it--while retaining our original levels of confidence. However, it is mysterious why, in cases in which we *do* have access to the original evidence, that evidence should play no role in determining what it is reasonable for us to believe but is rather completely swamped by the opinions that we form in response to it. It is a weakness of The Equal Weight View that it assimilates cases in which one does have access to the original evidence to cases in which one does not.

I find the suggestion that the original evidence makes no difference at all once we respond to it a strange one. Of course, others might not share my sense of strangeness, and even those who do might very well be prepared to live with this consequence, given that other considerations might seem to tell strongly in favor of The Equal Weight View. For this reason, I want to press the point by offering four additional arguments. I offer the first two arguments in the spirit of plausibility considerations, designed to further bring out what I take to be the counterintuitiveness of the suggestion that the original

evidence gets completely swamped by psychological facts about how we respond to it. The third and fourth arguments are considerably more ambitious, inasmuch as they purport to show that there is something approaching absurdity in this idea.

(1) *A Comparison: Interpersonal and Intrapersonal Conflicts.* Compare the question of how it is rational to respond to interpersonal conflicts between the beliefs of different individuals with the question of how it is rational to respond to intrapersonal conflicts among one's own beliefs. Suppose that one suddenly realizes that two beliefs that one holds about some domain are inconsistent with one another. In such circumstances, one has a reason to revise one's beliefs. But how should one revise them? We can imagine a possible view according to which, whenever one is in such circumstances, one is rationally required to abandon *both* beliefs. This view about how to resolve intrapersonal conflicts is the closest analogue to The Equal Weight View. But such a view has little to recommend it. In some cases of intrapersonal conflict, the reasonable thing to do might be to abandon both beliefs until further evidence comes in. 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 one's total evidence but the other is not. A normative view about how it is reasonable to resolve inconsistencies among one's beliefs which completely abstracts away from facts about which beliefs are better supported by one's evidence, and which would have one treat one's prior beliefs on a par, regardless of how well- or ill-supported they are by one's total evidence, would not be an attractive one. But the features which make such a view unattractive are shared by The Equal Weight View.

(2) *Implausibly Easy Bootstrapping.*¹⁴ Consider

Case 5. You and I both accept The Equal Weight View as a matter of theory. Moreover, we scrupulously follow it as a matter of practice. At time t_0 , each of us has access to a substantial, fairly complicated body of evidence. On the whole this evidence tells against hypothesis H: given our evidence, the uniquely rational

¹⁴ The objection raised in this section is due, in all of its essential features, to Aaron Bronfman. I utilize it here with his permission.

credence for us to have in H is .3. However, as it happens, both of us badly mistake the import of this evidence: you give credence .7 to H while I give it .9. At time t1, we meet and compare notes. Because we both accept The Equal Weight View, we converge on credence .8.

On The Equal Weight View, our high level of confidence that H is true at time t1 is automatically rational, despite the poor job that each of us has done in evaluating our original evidence. (Indeed, it would be unreasonable for us to be any less confident than we are at that point.) However, it is dubious that rational belief is so easy to come by. Suppose that when you and I meet to compare notes at time t1, I ask you for the evidence on the basis of which you invest such high credence in the hypothesis. You recite your evidence, evidence which in fact favors not-H over H. You then ask me for my evidence; I recite the same body of underwhelming considerations. According to The Equal Weight View, this process is sufficient to make it reasonable for both of us to have a high degree of confidence that H is true, despite the fact that, *ex hypothesi*, it was unreasonable for either of us to have a high degree of confidence before we met. But that seems mistaken.

(Suppose that another pair of peers has access to the same body of evidence. However, both of them evaluate the evidence correctly and thus give credence .3 to the hypothesis. Later they meet, compare notes, and thus maintain their original opinions. On The Equal Weight View, their later opinion about the hypothesis is no more reasonable than our later opinion is, for their opinion is the uniquely reasonable response to their total evidence, while our opinion is the uniquely reasonable response to our total evidence. If God is rewarding people in proportion to how well their beliefs reflect their evidence with respect to this particular hypothesis, He would have no basis for rewarding them more handsomely than He rewards us, since all four of us are doing maximally well. Again, that conclusion seems difficult to endorse.)

It is often noted that, at least on first inspection, The Equal Weight View would seem to have relatively radical implications for our actual practice.¹⁵ After all, many of us

¹⁵ Interestingly, this point is emphasized both by those who are sympathetic to The Equal Weight View as well as by those who seek to resist it. Examples of the former include Elga (forthcoming) and Feldman (2006); an example of the latter is van Inwagen (1996).

persist in retaining views that are explicitly rejected by those over whom we possess no discernible epistemic advantage. It seems that, if The Equal Weight View is true, then many of us should give up (or at least, become significantly less confident of) some of our deepest convictions about philosophy, politics, morality, history, religion, and other subjects in which there is substantial controversy among intelligent, thoughtful, and well-informed people. Indeed, advocates of The Equal Weight View sometimes devote substantial labor attempting to show that this *prima facie* consequence is not an *ultima facie* one.¹⁶ The operative idea, it seems, is that it is surely not so difficult for intelligent, thoughtful, and well-informed people to rationally hold confident beliefs about such matters. Although I reject The Equal Weight View, I myself do not think that this consequence (if it is indeed such) should be counted as a significant theoretical cost. On the contrary, the suggestion that many or most of us tend to be too confident of our controversial philosophical, political, historical (etc.) opinions strikes me as having considerable independent plausibility. What has thus far not been adequately appreciated about The Equal Weight View is to my mind a much more damning consequence. Namely, that if The Equal Weight View is true, then there will be cases in which rational belief is *too easy* to come by. That is, views for which there is in fact no good evidence or reason to think true can bootstrap their way into being rationally held simply because two irrationally overconfident peers encounter one another and confirm each other's previously baseless opinions.

Indeed, I believe that there is significantly worse trouble for The Equal Weight View on this front.

3.) *Even Easier, and More Implausible, Bootstrapping: Single person cases*

On The Equal Weight View, the evidence which determines what it is reasonable for us to believe in cases of peer disagreement consists in facts about the distribution of opinion among the peers. Let us call such evidence **psychological evidence**. Let us call the original evidence on which the peers base their opinions **non-psychological**

¹⁶ See especially Elga (forthcoming).

evidence.¹⁷ Above, we noted that there is at least one special case in which--as the advocate of The Equal Weight View would have it--it is highly plausible that what it is reasonable to believe is entirely fixed by the psychological evidence, viz. a case in which the psychological evidence is all the evidence that one has to go on. When one is aware of nothing relevant to some issue other than facts about the distribution of opinion, it is unsurprising that such facts suffice to fix what it is reasonable for one to believe about that question. In the even more special case in which one is aware of nothing relevant other than the distribution of opinion among a group of one's peers, one should give equal weight to each of their opinions. (Crucially, these thoughts are not the exclusive property of The Equal Weight View, a point to which we will return below.)

At one end of the spectrum then, are cases in which one's evidence is exhausted by psychological evidence concerning facts about the distribution of opinion (i.e., cases in which one's non-psychological evidence has dwindled to nothing). At the other end of the spectrum are cases in which all of one's evidence is non-psychological (i.e., cases in which one's psychological evidence has dwindled to nothing). Consider a case of the latter kind: at time t_0 , one possesses a body of non-psychological evidence E that bears on some question, but one is completely ignorant of what anyone else thinks about that question, nor has one yet formed an opinion about the issue oneself. Presumably, at this point a proponent of The Equal Weight View will agree that what it is reasonable for one to believe is wholly fixed by the non-psychological evidence (to the extent that what is reasonable to believe is fixed by the evidence at all). At time t_1 , one first forms an opinion about the hypothesis on the basis of this non-psychological evidence; let us suppose that one gives credence .7 to the hypothesis on the basis of the evidence. Assuming that one has access to facts about one's own confidence *via* introspection, one thus acquires one's first piece of psychological evidence that bears on the question. For

¹⁷ Some might find this terminology suboptimal on the grounds that *all* of one's evidence is ultimately psychological inasmuch as it consists of one's own psychological states. I think that this complaint rests on a mistaken view about the ontology of evidence, but no matter: one who thinks that all of our evidence ultimately consists of psychological states might read 'psychological evidence' and 'non-psychological evidence' as 'doxastic evidence' and 'non-doxastic evidence' in what follows.

one can now adopt a third person perspective on one's own opinion and treat the fact that one believes as one does as evidence which bears on the truth of the hypothesis. At time t_1 then, one's total evidence consists of one's original body of non-psychological evidence E , plus a single piece of psychological evidence, viz. the fact that one believes as one does. Call this new body of total evidence E_+ :

E_+ (one's evidence at time t_1)

The original body of non-psychological evidence E

The fact that one believes the hypothesis to degree .7

Suppose that at time t_2 one gains an additional piece of psychological evidence: one learns the opinion of a peer. Suppose that the peer gives credence .5 to the hypothesis. At time t_2 then, one's total evidence—call it E_{++} —consists of the following:

E_{++} (one's evidence at time t_2)

The original non-psychological evidence E

The fact that one believes the hypothesis to degree .7

The fact that one's peer believes the hypothesis to degree .5.

According to The Equal Weight View, one should split the difference with one's peer and believe the hypothesis to degree .6 at time t_2 ; we have criticized the view at some length on the grounds that it implausibly suggests that the psychological evidence swamps the non-psychological evidence in these circumstances. At present, however, I want to inquire about what a proponent of The Equal Weight View should say about what it is reasonable for one to believe back at time t_1 , when one knows one's own opinion about the hypothesis but no one else's. Does the psychological evidence swamp the non-psychological evidence *even then*? It would seem that the only principled answer for the proponent of The Equal Weight View to give to this question is 'Yes'. For the proponent of The Equal Weight View will insist that, at time t_2 , what it is reasonable for one to believe is determined by averaging the original opinions of the two peers; moreover, if, at an even later time t_3 , one becomes aware of the opinion of a third peer, then what it is

reasonable for one to believe will be determined by averaging the original opinions of the three peers; and if, at some still later time t_4 , one becomes aware of the opinion of a fourth peer... In general, for any time t_n , a proponent of The Equal Weight View will hold that what it is reasonable to believe is entirely fixed by averaging the opinions of the n peers. Why then should things be any different back at time t_1 , when the number of peers = 1? It seems as though the only principled, not *ad hoc* stand for the proponent of The Equal Weight View to take is to hold that the psychological evidence swamps the non-psychological evidence even when the psychological evidence is exhausted by what you yourself believe. On this view, before one forms some opinion about the hypothesis, how confident one should be that the hypothesis is true is determined by the non-psychological evidence; after one arrives at some level of confidence—in the present example, a degree of belief of .7—how confident one should be given the evidence that one then possesses is... .7. Of course, if one had responded to the original evidence in some alternative way—say, by giving credence .6 or .8 to the hypothesis--then the uniquely reasonable credence would be .6 or .8. On the picture of evidence suggested by The Equal Weight View, the distinction between believing and believing rationally collapses in cases in which one is aware of what one believes but unaware of what others believe. But that is absurd. Hence, The Equal Weight View is false.

Here we note an interesting general feature of The Equal Weight View and how it makes for trouble in the present case. On the operative conception of peerhood, peers resemble each other in possessing a similar general competence for assessing relevant evidence and arguments. If you regard someone as incompetent compared to yourself with respect to his or her ability to assess relevant considerations, then you do not regard that person as your peer. (As a relatively extreme case, we might think here of the relationship that the qualified teacher of philosophy stands in to those of her students who have not yet developed any sophistication in evaluating arguments.) Of course, in order to respond correctly to one's evidence on a given occasion, it is not sufficient that one is competent to do so; one must actually manifest one's competence. Even against a general background of competence, one might still over- or underestimate one's evidence on a given occasion: one commits a performance error, as it were. Notice that it is characteristic of The Equal Weight View to credit the views of others in proportion to

their general competence while abstracting away from facts about actual performance. What it is reasonable to believe in cases of peer disagreement is in effect determined by taking the average of peer opinion; crucially, in this calculation, the opinions that have been arrived at *via* the commission of performance errors will count for just as much as those opinions that are appropriate responses to the shared evidence.¹⁸ Bare truths about who has in fact manifested their underlying competence and who has not make no difference in cases of peer disagreement. However, once facts about general competence are privileged in this way in multi-person cases, it seems arbitrary and unmotivated to continue to maintain that actual performance makes a significant difference in single person cases (i.e., cases in which a single individual arrives at an opinion on the basis of the non-psychological evidence that he possesses). Rather, on the suggested picture, if I am generally competent in the way that I respond to evidence (and I know that I am), then this should be enough to guarantee that I am reasonable in responding to my evidence in whatever way that I do. But this contradicts our initial assumption, viz. that one way of ending up with an unreasonable belief is to respond incorrectly to one's evidence, despite possessing the ability to respond to that evidence correctly.

4.) *The Litmus Paper Objection*

Let us set aside, for the moment, the special case of disagreement among peers, and reflect on a much more general question: in what circumstances does it make sense for

¹⁸ At least, so long as one has no *independent* grounds for attributing such performance errors. Of course, it is open to a proponent of The Equal Weight View to say that, even if you and I possess similar general competence, it is permissible for you to discount my opinion when (e.g.) you notice that I was distracted while surveying the evidence in a way that you were not, or that I did so while under the influence of some temporarily mind-numbing drug, or so on. What the proponent of The Equal Weight View will *not* allow is that my actually having committed a performance error can make a difference when your only grounds for attributing such an error to me consists in the fact that I have arrived at (what you take to be) an incorrect answer to the question about which we disagree. It is this feature of The Equal Weight View which distinguishes it from the alternative view that I will offer and which leaves it vulnerable to the current objection.

me to treat the fact that someone else believes as she does as evidence for the truth of that which she believes? A true (although perhaps not especially informative) answer: exactly when I take her belief to be a reliable indication of how things stand in the relevant part of reality. Thus, suppose that I know, on the basis of extensive past experience, that when my weather forecaster judges that it will rain the next day, it tends to rain 80% of the time. In that case, I will treat her judgments to the effect that it will rain as evidence that it will rain, inasmuch as I take there to be a positive correlation between the two. Notice that, in this respect, there is absolutely nothing special about the way in which the judgments of another person come to count as evidence. Compare: I treat the fact that the litmus paper turns red as evidence that the liquid in which it is immersed is an acid because, on the theories that I accept, the former is a reliable indication of the latter. This seems perfectly parallel to the reason why I treat the fact that my weather forecaster expects it to rain tomorrow as evidence that it will rain tomorrow. In general, the way in which the judgments of some other mind come to play the role of evidence does not differ from the way in which other states of the world do.

I believe that this observation, while elementary, is already enough to cast significant doubt on the Equal Weight View. For consider your perspective, as one attempting to determine what to believe about some proposition. You carefully survey what you take to be your evidence: various states of the world, the obtaining of which you take to provide clues as to whether the proposition is true or false. Some of these states of the world are bits of psychological reality, the beliefs of others—that Smith is highly confident that the proposition is true, that Jones is less so, and so on. Others of these states of the world are bits of non-psychological reality—for example, the fact that the litmus paper turned a given color in such-and-such circumstances. Insofar as you think it relatively unlikely that some part of psychological reality would be as it is unless the proposition were true, you regard the fact that things are arranged thus and so as evidence that speaks in favor of the proposition. But by the same token, insofar as you think it relatively unlikely that some piece of non-psychological reality would be as it is unless the proposition were true, you regard the fact that things are arranged that way as evidence that speaks in favor of the proposition. Now consider the special case in which you possess a considerable amount of non-psychological evidence, but where your

psychological evidence is exhausted by the fact that (i) you yourself believe the proposition to degree n , and (ii) some peer believes the proposition to degree m , where m differs from n . Again, on The Equal Weight View, what it is reasonable for you to believe is the average of m and n ; in effect, on The Equal Weight View, one ought to give no weight to the non-psychological evidence in the presence of the psychological evidence. But what could be the rationale for such a policy of invidious discrimination? Why should the psychological evidence count for everything, and the non-psychological evidence for nothing, *given that the way in which the two kinds of evidence qualify as such is exactly the same?*

The idea that the non-psychological evidence is wholly swamped by the psychological evidence will look increasingly implausible the more the former is made to resemble the latter. Suppose that, in addition to my keen interest in having accurate beliefs about the future states of the weather, I am also a skilled engineer; I thus set out to construct a machine that will be of use in weather forecasting. My intention in constructing the machine is that its states will track future states of the weather: ideally, the machine will be in State A on a given day if and only if it will rain at some point during the following day, and so on. Unfortunately, the machine turns out to be a highly imperfect indicator of impending rain: I know, on the basis of significant past experience, that 60% of the time that the machine enters into State A, it rains the next day, while 40% of the time it does not. Suppose further that this is roughly how reliable *I* am when it comes to predicting the rain, and therefore, how reliable my peers are (and that I am aware of this). Now consider a case in which I have access to a large and complicated body of non-psychological evidence which bears on the question of whether it will rain tomorrow; one element of this body is the fact that the machine is currently in State A. At time t_0 , I survey this evidence and arrive at an opinion about whether it will rain tomorrow; at the same time, a peer does the same; subsequently, at time t_1 , I learn of my peer's opinion. Now let us add the following to the story. Coincidentally, both my peer and I failed to give any weight to the fact that the machine was in State A in arriving at our respective opinions (we can imagine that although this relevant fact was known to both of us, it was simply not factored in amidst the large mass of relevant data). The fact that the machine was in State A thus plays no role in influencing what either of us believes at time t_0 , and

therefore, according to The Equal Weight View, makes no difference to what it is reasonable for me to believe at time t_1 . On the other hand, the fact that my peer believes as she does makes a great deal of difference to what it is reasonable for me to believe at that point. Notice how differently the advocate of The Equal Weight View will treat the opinion of my peer and the state of the machine, despite their evident structural similarities. If I proceed to give no weight to the opinion of my peer at time t_1 , and instead simply retain my original opinion utterly uninfluenced by the fact that he believes as he does, then my behavior is a paradigm of unreasonableness. However, when at the same time I give no weight to the fact that the machine is in state A, my belief can still be perfectly reasonable. But this is odd, given that the way in which my peer's opinion is related to future states of the weather and the way in which the state of the machine is related to future states of the weather would seem to be very much the same.¹⁹

Perhaps The Equal Weight Theorist will insist instead that I should treat the machine as a peer with respect to weather forecasting. On this line of thought, in building the machine I am in effect constructing a robot peer; the machine's believing that it will rain tomorrow consists in its being in State A. However, this maneuver seems somewhat desperate and will not always be available. (The fact that the litmus paper turns red does not constitute its coming to believe that the liquid in which it is immersed is an acid.)

Although I will return for a few parting shot below, this concludes my brief against The Equal Weight View. Even if one thinks that there is some merit to the arguments of this section, however, one might still be reluctant to abandon The Equal Weight View. For one might think that it is strongly supported by certain kinds of cases, e.g., cases involving conflicting perceptual judgments. I will address such concerns directly in Section 5 below. To briefly anticipate: while I agree that The Equal Weight View returns the correct verdicts about such cases, I do not believe that this fact is ultimately of much dialectical significance. The reason for this is straightforward: what I take to be the best

¹⁹We might also imagine cases in which I am confused, or it is temporarily opaque to me, whether a given piece of evidence that favors the proposition that it will rain tomorrow consists of (i) the fact that my peer believes that it will rain tomorrow, or, alternatively, (ii) the fact that the machine is in State A. Will the proponent of The Equal Weight View insist that, once I learn the truth, a significant revision in my opinion about whether it will rain tomorrow might be in order?

alternative to The Equal Weight View yields the same verdicts when applied to the cases in question. Before explicitly addressing those considerations which seem to favor The Equal Weight View then, I will put my preferred alternative on the table.

4. The Total Evidence View

Against The Equal Weight View, I have argued that there is at least one type of situation in which one is not required to split the difference with a peer who disagrees. Specifically, if you and I have arrived at our opinions in response to a substantial body of evidence, and your opinion is a reasonable response to that evidence while mine is not, then you are not required to give equal weight to my opinion and to your own. Indeed, one might wonder whether you are required to give *any* weight to my opinion in such circumstances. Suppose that when we first discover that we hold different opinions about the issue in question, neither one of us is aware of why the other believes as he does. Naturally enough, you inquire as to why I believe as I do: because you regard me as in general a competent evaluator of evidence, you suspect that the fact that we believe differently about the issue is underwritten by my having access to some crucial piece of evidence to which you lack access (or vice versa). However, this suspicion turns out to be disappointed: I myself insist that I have no additional evidence, apart from considerations of which you yourself were already aware. In arriving at your opinion about the issue, you have thus taken into account every consideration which I would cite as a reason to justify my own opinion. Given this, why would you be rationally required to give some additional weight to my opinion (an opinion which is in fact unreasonable) and move from your own, an opinion which is in fact a perfectly reasonable response to the only evidence that either one of us claims to have?

Recall from above

The No Independent Weight View: In some cases of peer disagreement, one might be perfectly reasonable even if one gives no weight at all to the opinion of one's peer.

and

The Symmetrical No Independent Weight View: In some cases of peer disagreement, both parties to the dispute might be perfectly reasonable even if neither gives any weight at all to the opinion of the other party.

In section 2, I tentatively suggested that these views might be defended by appeal to the apparent possibility of ‘reasonable disagreements’: cases in which there is a range of rationally permissible attitudes to take towards some proposition given one’s evidence. However, for the sake of argument, we are currently supposing that The Uniqueness Thesis is true, and thus, that the apparent possibility of reasonable disagreements is merely apparent: in fact, whenever two individuals believe different things on the basis of the same evidence, at least one of the two is being unreasonable, or at least, less reasonable than the other. Assuming that The Uniqueness Thesis is true, The Symmetrical No Independent Weight View is false. However, even if The Symmetrical No Independent Weight View is false, The No Independent Weight View might still be true. For even if it cannot be reasonable for both you and I to give no weight to the other’s opinion, perhaps it is nevertheless reasonable for you to give no weight to my opinion in the envisaged circumstances. As formulated above, The No Independent Weight View states that it might be perfectly reasonable to give no weight to the opinion of one’s peer ‘in some cases’. We have now arrived at a proposal for what the relevant class of cases is, viz. the class of cases in which one’s original opinion correctly reflects the evidence that one shares with one’s peer but his opinion does not. Consider then

The Asymmetrical No Independent Weight View: In cases of peer disagreement, it is reasonable to give no weight to the opinion of a peer as long as one’s own opinion is the reasonable response to the original evidence.

On this view, if either of the two peers engaged in a disagreement has in fact evaluated their shared evidence correctly, then that peer should stick to her guns, and the other peer should convert, since the opinion in question is the one which is in fact best supported by their evidence.

However, The Asymmetrical No Independent Weight View is false. Even if one responds to the original evidence in an impeccable manner and one’s peer does not, the fact that one’s peer responds as he does will typically make it rationally incumbent upon

one to move at least some way in his direction. First let us satisfy ourselves that this is so; we will then inquire as to why it is so.

Consider

Case 6. You are a professional mathematician. Within the mathematics community, there is substantial and longstanding interest in a certain mathematical conjecture. (Call it **The Conjecture**.) If forced to guess, some members of the community would guess that The Conjecture is true, others that it is false; all agree that there is no basis that would justify a firm opinion one way or the other. Then, one day, the unexpected happens: alone in your study, you succeed in proving The Conjecture. On the basis of your proof, you become extremely confident, indeed practically certain, that The Conjecture is true. Because your high degree of confidence is based on a genuine proof that you correctly recognize as such, it is fully justified. Later, you show the proof to a colleague whose judgment you respect. Much to your surprise, the colleague, after examining the proof with great care, declares that it is unsound. Subsequently, you show the proof to another colleague, and then to a third, and then to a fourth. You approach the colleagues independently and take pains to ensure that they are not influenced by one another in arriving at their judgments about the status of your proof. In each case, however, the judgment is the same: the proof is unsound. Ultimately, your proof convinces no one: the entire mathematical community is united in its conviction that it is unsound, and thus, that the status of The Conjecture remains very much an open question.

In the face of this consensus, it would be unreasonable for you to remain practically certain that The Conjecture is true. You should be less confident of The Conjecture after your proof has been deemed unsound by the mathematical community than you were immediately after you first proved The Conjecture, back when you were alone in your study. Of course, because the proof is in fact sound, the judgment of the community to the contrary is misleading evidence, evidence that points in the wrong direction. But misleading evidence is evidence nonetheless, and the acquisition of such evidence will typically make a difference to what it is reasonable for one to believe. Moreover, if you are rationally required to be less confident after *all* of your peers have disagreed with you, then it would seem that you are also required to be at least somewhat less confident after even *one* of your peers disagrees with you. For suppose that it was rationally permissible to give zero weight to the opinion of the first colleague. In that case, you could have left her office as rationally confident as when you entered, in which cases you would have been in the same state of practical certainty upon entering the office of the

second colleague that you consulted. Indeed, in that case it seems that you might as well simply forget about the fact that the whole unpleasant business with the first colleague occurred at all before visiting the second colleague, in which case you would be in more or less exactly the same position upon entering the office of the second colleague. And if it is rationally permissible to give zero weight to *his* opinion...

Moral: the fact that a peer believes differently can make it rationally incumbent upon you to change what you currently believe, even if, had the peer responded to the evidence in a reasonable manner, he too would believe exactly as you believe. One should give some weight to one's peer's opinion even when from the God's eye point of view one has evaluated the evidence correctly and he has not. But why? Exactly because one does not occupy the God's eye point of view with respect to the question of who has evaluated the evidence correctly and who has not.²⁰ Typically, when one responds reasonably to a body of evidence, one is not utterly blind to the fact that one has done so; on the other hand, such facts are not perfectly transparent either. Even if one has in fact responded to the evidence impeccably on a given occasion, one might still have reason to doubt that one's performance was impeccable. Such a reason is provided when a peer responds to that same evidence differently. To give no weight to the fact that a peer responds to the evidence differently is in effect to treat it as certain that one's peer is the one who has misjudged the evidence. But it would be unreasonable to be certain of this, even when it is true.²¹

²⁰ Cf. the lucid and illuminating discussion of this point in Christensen (2007) and (forthcoming).

²¹ In Kelly (2005), I suggested that we should regard the views of a generally reasonable person as in effect providing *higher order evidence*: that is, evidence about the normative upshot of the evidence to which she has been exposed. (See especially the discussion on pages 185-190). So, for example, the fact that a generally reasonable person *S* believes *p* is (defeasible) evidence in favor of the epistemic proposition that *it is reasonable to believe p given S's evidence*. I emphasized that higher order evidence of this sort bears most directly on epistemic propositions and that acquiring such evidence will often make a straightforward difference to what it is reasonable for one to believe *about* particular bodies of evidence. On the other hand, I expressed considerable skepticism about the idea that the higher order evidence provided by the fact that a generally reasonable person believes a given proposition will also make a difference to what it is reasonable for one to believe about that proposition in a case in which one knows that one already possesses all

Rationality consists in responding appropriately to one's evidence. But one's evidence includes evidence to the effect that one does not always respond appropriately to one's evidence (i.e., evidence to the effect that one is fallible in responding appropriately to one's evidence), as well as evidence to the effect that one is more likely to have responded inappropriately when one finds oneself in certain circumstances. When one possesses higher order evidence to the effect that one is currently in circumstances in which one is more likely than usual to have made a mistake in responding to one's first order evidence, one has a reason to temper one's confidence—even if that confidence is in fact an impeccable response to the first order evidence. When one finds oneself in the position of a minority of one in the way that one has responded to the evidence, one should temper one's confidence, for one now possesses higher order evidence that suggests that the bearing of the original, first order evidence is something other than what one initially took it to be. Moreover, this is so even if the higher order evidence is misleading, as when one has in fact responded appropriately to the first order evidence and one's peers have not.

The tendency of higher order considerations to affect how much confidence one should invest in the deliverances of one's first order reasoning is a quite general phenomenon; it will be helpful to briefly explore the general phenomenon in order to gain some perspective on the case at hand. Consider first a case of practical deliberation: one is attempting to determine which of two attractive job offers to accept. Some considerations point in one direction; other considerations point in the opposite direction. One deliberates and concludes that, on balance, one's practical reasons favor accepting one of the two offers over the other. How confident should one be that this is what one should do? This can depend, not only on the strength of the reasons for and against accepting that offer but also on (what one knows about) the circumstances in which one is deliberating. If one knows that one is deliberating in circumstances in which one's

of the evidence on which the person bases her belief. (Foremost among my reasons for skepticism: the 'double-counting' argument rehearsed on pages 187-188.) What I say here constitutes a departure from the earlier skeptical attitude: on the present view, higher order evidence about the bearing of one's first order evidence *is* typically relevant to what it is reasonable to believe on the basis of that evidence.

ability to weigh practical reasons tends to be compromised (e.g., one is in a state of inebriation), then one should be less confident than if one is deliberating in more ideal circumstances. Moreover, this is so even if one has in fact deliberated impeccably, despite being inebriated, and has assigned the correct weights to all of the practical reasons which bear on the choice. Of course, higher order considerations having to do with whether one is deliberating in the bad case (one is inebriated) or in the good case (one is sober) are not themselves reasons for or against accepting either job offer in the way that the practical reasons which bear directly on the choice are; nevertheless, they too make a difference to what it is appropriate to conclude on the basis of one's deliberations. On the present picture, even if one's practical reasoning in the good case is identical to one's practical reasoning in the bad case, what one is justified in concluding on the basis of that reasoning might differ.

The same point holds for theoretical reasoning. I am entitled to have more confidence in the conclusion of a given piece of mathematical reasoning when I have performed the relevant calculation in the morning when I am wide awake, than if I arrive at the same conclusion by reasoning in exactly the same way late in the evening, when I know that I am prone to making mistakes due to fatigue.

On the present view, cases in which one in fact responds impeccably to one's evidence, but one's peer responds inappropriately, are much like cases in which one engages in a flawless piece of practical reasoning despite being inebriated. The fact that a peer has responded to the evidence differently should lead one to temper one's confidence in one's own response, just as the fact that one is inebriated should lead one to temper one's confidence in the conclusion of one's practical reasoning despite the actual flawlessness of one's performance. Again, in both cases, it is the fact that the status of one's performance is not perfectly transparent that opens the door for higher-order considerations to make a difference.

Of course, to acknowledge that higher order considerations make *some* difference is not to fall back into the mistake of thinking that they make *all* the difference. After all, even when one's current level of inebriation makes it significantly more likely that one will over- or underestimate the strength of one's practical reasons (and one knows that this is so), one can still make more or less rational decisions, and the status of a given

decision will typically depend a great deal on the overall disposition of those practical reasons. Similarly for the theoretical case: although you should be somewhat less confident that The Conjecture is true upon finding that a colleague remains unconvinced despite having been presented with your proof, it is a mistake to think that at that point the only evidence that makes a difference are the respective psychological reactions of you and your colleague. When one possesses what is in fact a genuine proof that one correctly recognizes as such, one possesses an extremely strong piece of evidence. (Indeed, it would perhaps be difficult to imagine a stronger single piece of evidence for anything.) The justification afforded by such a piece of evidence has a certain robustness in the face of challenge: it is not easily washed away by the fact that another mistakenly fails to appreciate it on a given occasion. Of course, your colleague might feel just as confident that your proof is unsound as you feel that it is sound. Indeed, all of the psychological accompaniments of the two judgments might be the same. But in any case, we have independent reason to be skeptical of the idea that phenomenology is that on which epistemic status supervenes. In general, when one reasons badly, one's phenomenology might be indistinguishable from one's phenomenology when one reasons impeccably (in both cases, one has the same feelings of subjective certainty, and so on). We should not thereby be driven to the conclusion that the deliverances of good reasoning and bad reasoning have the same epistemic status.²²

Where does this leave us?

In the last section, I argued that, in cases of peer disagreement, getting the original, first order evidence right typically counts for *something* (pace The Equal Weight View). In this section, I have argued that doing so does not count for *everything* (pace The No Independent Weight View). Indeed, from the present perspective, there is a sense in which The Equal Weight View and The No Independent Weight View both suffer from the same fault: they embody overly simple models of how one's first order evidence and one's higher order evidence interact in determining facts about what it is reasonable to believe all things considered. On The Equal Weight View, what it is reasonable to

²²Recent--and to my mind, compelling--critiques of the idea that there is any interesting and important epistemic status which supervenes on phenomenology are provided by Timothy Williamson (2000) and Ernest Sosa (1999, 2002, 2007)

believe in cases of peer disagreement in effect supervenes on facts about the distribution of peer opinion. On The No Independent Weight View, what it is reasonable to believe in such cases supervenes on facts about the first order evidence possessed by the peers. On the present view, both of these supervenience claims are false: neither class of facts suffices on its own to fix the facts about what it is reasonable to believe. Rather, what it is reasonable to believe depends on both the original, first-order evidence as well as on the higher-order evidence that is afforded by the fact that one's peers believe as they do. For this reason, it seems appropriate to call the view on offer **The Total Evidence View**.

Even if both The Equal Weight View and The No Independent Weight View are unsatisfactory, we might still wonder: which is closer to the truth? Granted that on The Total Evidence View both the first order evidence and the higher order evidence count for something, which kind of evidence plays a greater role in fixing facts about what it is reasonable to believe?

It is a mistake, I believe, to think that there is some general answer to this question. In some cases, the first order evidence might be extremely substantial compared to the higher-order evidence; in such cases, the former tends to swamp the latter. In other cases, the first order evidence might be quite insubstantial compared to the higher order evidence; in such cases, the latter tends to swamp the former. (We will consider plausible examples of each of these types of case below.) In still other cases, the two kinds of evidence might play a more or less equal role in fixing facts about what it is reasonable to believe. So the question of which counts for *more*—peer opinion, or the evidence on which the peers base their opinion?—is not, I think, a good question when it is posed at such a high level of abstraction.

Nevertheless, we can offer some general observations which bear on this issue here. Consider again the kind of case that we have employed in attempting to undermine The Equal Weight View: initially, you and I have access to the same substantial body of evidence E, evidence which in fact strongly favors H over not-H; you respond reasonably and so are quite confident that H is true; I on the other hand respond unreasonably and am equally confident that H is false. Once we compare notes, our new total evidence consists of E*:

- (i) Our original evidence E
- (ii) The fact that you are quite confident that H is true
- (iii) The fact that I am quite confident that H is false

What is it reasonable for us to believe about H on total evidence E*? Given that you and I are peers, it is plausible to suppose that the two pieces of higher order psychological evidence ((ii) and (iii)) are more or less equally strong pieces of evidence that point in opposite directions. All else being equal then, one would expect E* to favor H over not-H inasmuch as it is composed of a substantial body of evidence that strongly favors H over not-H, supplemented by two additional pieces of evidence of approximately equal strength, one of which tends to confirm H, the other of which tends to disconfirm H.

Indeed, it is tempting to think that, if in fact our respective psychological reactions count as more or less equally strong pieces of evidence that point in opposite directions, then they in effect cancel each other out and leave what it is reasonable for us to believe unchanged. According to this line of thought, what it is reasonable for us to believe about H on E* is identical to whatever it was reasonable for us to believe about H on E, inasmuch as the net effect of adding the two new pieces of evidence comes to zero. Here The Asymmetrical No Independent Weight View threatens to return *via* the back door, at least in a special class of cases, viz. cases in which peer opinion is evenly divided. For in such cases, the evidence afforded by peer opinion is perfectly counterbalanced.

However, this tempting line of thought is mistaken. The addition of the counterbalanced psychological evidence *does* make a difference to what it is reasonable for us to believe. For once the counterbalanced evidence is added to our original evidence, a greater proportion of our total evidence supports an attitude of agnosticism than was previously the case; the evidence available to us now is on the whole less supportive of H than before. The addition of (ii) and (iii) thus has a moderating impact and tends to push what it is reasonable for us to believe about the hypothesis in the direction of agnosticism. Therefore, given that E is a substantial body of evidence that strongly favors H over not-H, we would expect that E* will also favor H over not-H, although not to as a great a degree as E does. (That is, all else being equal, the reasonable

level of confidence to have in hypothesis H on evidence E^* will be greater than .5 but less than whatever it was reasonable to have on evidence E.)

Significantly, the point generalizes beyond the two person case. As more and more peers weigh in on a given issue, the proportion of the total evidence which consists of higher order psychological evidence increases, and the proportion of the total evidence which consists of first order evidence decreases. As the number of peers increases, peer opinion counts for progressively more in determining what it is reasonable for the peers to believe, and first order considerations count for less and less. At some point, when the number of peers grows large enough, the higher order psychological evidence will swamp the first order evidence into virtual insignificance. In such cases, The Total Evidence View becomes more or less extensionally equivalent to The Equal Weight View with respect to what it requires the peers to believe. Moreover, this holds regardless of the particular way in which opinion is distributed among the peers. That is, it holds for cases in which peer opinion is evenly divided, for cases in which peer opinion is unanimous, as well as for intermediate cases.

Imagine an infinite number of peers confronted with a finite amount of evidence which bears on some issue. Each of the peers inspects the evidence and independently arrives at a view. When the peers compare notes, they find that opinion among them is perfectly divided: every peer on one side of the issue has one and only one counterpart on the other side. In these circumstances, the peers should suspend judgment about the issue, even if that response is not the most rational response to the original, first-order evidence. With respect to this case, The Equal Weight View returns the correct verdict from the perspective of one who holds The Total Evidence View. This is so *not* because the higher order evidence trumps the first order evidence in general, as the proponent of The Equal Weight View maintains. Rather, it is because in sufficiently extreme cases, the higher order psychological evidence might be so substantial compared to the first order non-psychological evidence that the former in effect swamps the latter into virtual insignificance.

The same holds true for cases in which the peers find that they agree. Above, we looked askance at the idea that two peers, both of whom irrationally hold some view that is not in fact supported by their evidence, might bootstrap their way into rationally

holding that view simply by encountering one another and comparing notes. Indeed, we took the fact that The Equal Weight View licenses such two-person bootstrapping as a consideration that counts against it (See 3.2 above). However, as the number of generally reliable peers who independently respond to their evidence in the same mistaken manner increases, such bootstrapping seems less and less objectionable. At some point, it becomes, I believe, unobjectionable. If I hold some belief on the basis of fallacious reasoning, then it will typically not be reasonable for me to hold that belief. However, in the unlikely but possible situation in which a large number of generally reliable peers mistakenly arrives at the same conclusion by independently committing the same fallacy, it will typically be reasonable for them to believe that conclusion upon comparing notes, even if there is no legitimate first-order reasoning by which they could have arrived at the conclusion. Again, in this case The Equal Weight View yields the correct verdict from the perspective of The Total Evidence View. As before, this is not due to some general tendency of higher order evidence to trump first order evidence. Rather, it is due to the fact that in this case, the higher order evidence which has been amassed is sufficiently substantial compared to the first order evidence that it effectively determines the bearing of the overall evidence.

Does this in effect give the game away to someone who takes the diversity of opinion with respect to various controversial issues to mandate an attitude of agnosticism about those issues? That is, even if The Equal Weight View is false and The Total Evidence View is true, won't all of the interesting/threatening/radical consequences that seemed to follow from The Equal Weight View still be true, at least if one is sufficiently generous in attributing the status of 'peer' to other people? Isn't agnosticism the only reasonable stance to take towards all of those controversial issues on which peer opinion is heavily divided, as the proponent of The Equal Weight View has insisted all along?

Consider also those philosophical questions with respect to which there is consensus, or near consensus. Suppose, plausibly, that there are very few if any genuine skeptics about other minds: informed philosophical opinion is (close to) unanimous in holding that one is typically in a position to know that there are minds other than one's own. In Kelly (2005), I took a dim view of the suggestion that this fact would make it unreasonable to embrace skepticism about other minds: rather, whether it is reasonable or unreasonable to

embrace skepticism about other minds is primarily a matter of the quality of the first-order arguments for and against such skepticism, arguments that do not make reference to empirical, sociological facts about the number of skeptics and non-skeptics. However, in light of the present view, a reversal of this judgment might seem to be in order. Could it really be that the unreasonableness of skepticism about other minds consists in the *unpopularity* of such skepticism among the relevant class of people?

Before acquiescing in this line of thought, we should note an important element of idealization in our discussion to this point, an element which looms large in the present context. Throughout, we have been concerned with the probative force of peer opinion in cases in which the peers arrive at their opinions *independently* of one another. This assumption of independence tends to maximize the probative force of peer opinion relative to the probative force of first-order evidence. Impressive evidence that a given answer to a question is the correct answer is afforded when a large number of generally reliable peers independently converge on that answer. On the other hand, the less their convergence is an independent matter, the less weight such convergence possesses as evidence.²³ Similarly, evidence that strongly favored agnosticism with respect to some question would be a more or less even distribution of opinion among a substantial number of peers, where each of the peers has arrived at his or her own opinion independently of the others. Again, the less such independence is present, the weaker the higher order evidence will be relative to the first-order evidence.

Consider, as an especially extreme illustration of the importance of independence, the venerable ‘Common Consent’ Argument for the Existence of God. In its simplest and most straightforward form, the argument runs as follows:

(Premise) Everyone believes that God exists.

(Conclusion) Therefore, God exists.

²³ On the importance and nature of independence, see especially the illuminating discussion in Goldman (2001), pp.150-156. In that paper Goldman is specifically concerned with the interesting question of how a non-expert should respond to disagreement among the experts, but the analysis of independence which he offers would seem to be highly relevant to a host of other important issues in social epistemology as well.

(In a slightly less crude form, the premise of the argument is that *almost* everyone, or the great majority of humankind, believes that God exists.²⁴)

As arguments go, The Common Consent Argument for the Existence of God is not exactly an overwhelming one, possessing as it does the twin defects of transparent invalidity and the having of an obviously false claim as its sole premise. Nevertheless, even though *God exists* does not follow from *Everyone believes that God exists*, we can ask: if it were true that everyone, or almost everyone, believed that God exists, how much support would that lend (if any) to the proposition that God exists?

This is a complicated question about which much could be said; here we note the following. Whatever evidence is afforded for a given claim by the fact that several billion people confidently believe that that claim is true, that evidence is less impressive to the extent that the individuals in question have not arrived at that belief independently. That is, the evidence provided by the fact that a large number of individuals hold a belief in common is weaker to the extent that the individuals who share that belief do so because they have influenced one another, or because they have been influenced by common sources. (I assume that both of these conditions play a large role in the case of religious belief.) In principle, the fact that a small handful of people arrive at the same belief independently of one another might be better evidence that that belief is true, than if many millions of people arrive at the same belief non-independently. The intellectual case for Islam would not be any stronger today if birthrates in Muslim countries had been twice as high in past decades as they actually were; nor would the case be any weaker if such birthrates had been significantly lower.

²⁴ Perhaps unsurprisingly, the Common Consent Argument is not taken very seriously any more, even in those circles in which arguments for the existence of God are still taken seriously. It is, for example, rarely if ever included among the usual rogue's gallery of arguments for the existence of God (the ontological argument, the cosmological argument, etc.) in anthologies or course syllabi devoted to the philosophy of religion. Historically, however, it was taken *quite* seriously. A list of prominent thinkers who endorsed some recognizable variant of it would include Cicero, Seneca, the Cambridge Platonists, Gassendi, and Grotius; in addition, it was discussed critically by (among many others) both Locke and Mill. For an overview, see the useful survey in Edwards (1967).

The same holds for cases in which there is widespread disagreement but where the members of the contending factions have not arrived at their opinions independently. In an interesting recent essay, G.A. Cohen notes that the Oxford trained philosophers of his generation are almost unanimously of the opinion that there is a philosophically important distinction between analytic and synthetic truths. But on the other hand

...people of my generation who studied philosophy at Harvard rather than at Oxford for the most part *reject* the analytic/synthetic distinction. And I can't believe that this is an accident. That is, I can't believe that Harvard just *happened* to be a place where both its leading thinker rejected that distinction and its graduate students, for independent reasons—merely, for example, in the independent light of reason itself—also came to reject it. And vice versa, of course, for Oxford. I believe, rather, that in each case students were especially impressed by the reasons respectively for and against believing in the distinction, because in each case the reasons came with all the added persuasiveness of personal presentation, personal relationship, and so forth (p.18, *emphases his*).

Consider Cohen's position as one attempting to determine what to believe about this issue. On the one hand, there are the first order considerations that have been offered for and against the existence of a philosophically significant analytic-synthetic distinction. In addition, Cohen is also aware of the views of other individuals who are similarly acquainted with those first order considerations and whom he regards as his peers in other relevant respects. In weighing evidence of the latter kind, Cohen should sharply discount for the fact that (as he sees it) many individuals on both sides of the issue hold the views that they do because those views were held by their teachers. That is, in the counterfactual situation in which the distribution of peer opinion is exactly as it is, but in which each of the peers arrived at his or her view in response to 'the independent light of reason itself', the higher order evidence possessed by Cohen would be much more substantial than it is as things actually stand. The point is not that individuals who believe what their teachers believe are less reliable than they would be if they made up their own minds. Indeed, as a general matter, this is not even true. (If your teacher is better at assessing the arguments than you are, then you will be more reliable if you simply believe as she does than if you arrive at a view on the basis of your own assessment of the arguments.) The point, rather, is that insofar as one believes as one

does because this is what one's teacher believes, the fact that one believes as one does is not an *additional* piece of psychological evidence, over and above the psychological evidence afforded by the teacher's belief.

The general moral: even in cases in which opinion is sharply divided among a large number of generally reliable individuals, it would be a mistake to be impressed by the sheer number of such individuals on both sides of the issue. For numbers mean little in the absence of independence. (It is, of course, an empirical question—one that belongs, presumably, to psychology and sociology-- how independently people arrive at their views about various issues.) If one uncritically assumes that the members of the contending factions have arrived at their views independently, then one will tend to overestimate the importance of other people's opinions as evidence and underestimate the importance of the first order evidence and arguments. One will be too quick to conclude that agnosticism is the reasonable stance in cases in which opinion is sharply divided, and too quick to conclude that deference to the majority is the reasonable course in cases in which opinion is not sharply divided.²⁵

Nevertheless, it is true that on The Total Evidence View, there will be possible cases in which the higher order evidence is sufficiently substantial compared to the first order evidence that the latter counts for (almost) nothing. By the same token, however, there will be possible cases in which the opposite is true. What is a case in which peer opinion effectively counts for nothing in virtue of being overwhelmed by the first order considerations? Consider a case discussed by both Christensen (2007, pp.199-203) and Elga (forthcoming, p.xxx). You and I go to dinner with several friends; at the end of the meal we independently calculate what an individual share of the total bill comes to (imagine that the group has agreed to split the bill evenly among its members). You judge that an individual share is \$43 per person, a perfectly plausible (and, let us suppose, correct) answer to the question of what each of us owes. I, however, arrive at an absurd

²⁵ Indeed, as Hartry Field pointed out to me, the need to discount the numbers is not limited to cases in which there is causal dependence present, as in the examples considered above. If I know that two individuals will respond to given evidence in the same manner, then I should treat their having arrived at some particular answer as *one* piece of evidence, and not two pieces of evidence, in favor of that answer (even if their both having arrived at that answer is in no way underwritten by some causal link).

answer of \$450, an amount which significantly surpasses the total bill. Both Christensen and Elga think that, in *these* circumstances, you are *not* required to treat my answer and your answer with equal respect; indeed they think that you are entitled to more or less dismiss my answer entirely. The difficulty is how to account for this on a picture according to which splitting the difference is typically the appropriate response to peer disagreement. In general, it is at least a *prima facie* embarrassment for The Equal Weight View that the following is possible: a person for whom one has arbitrarily strong evidence that he or she is a peer might nevertheless give a patently absurd answer on a given occasion. For it seems incredible that, in such circumstances, one would be unreasonable if one failed to treat the peer's patently absurd answer and one's own non-absurd answer even-handedly.

Unsurprisingly, both Christensen and Elga have interesting and detailed stories to say about why, in these but not in otherwise similar cases, one need not give any weight to the view of one's peer.²⁶ We will not pause to evaluate the specifics of their respective proposals; here we note only how The Total Evidence View offers an extremely straightforward and compelling explanation of why you are entitled to effectively discount my absurd opinion. Quite simply: given the totality of considerations available to you which bear on the question at issue (e.g., your knowledge that the total bill is n , a number which is less than \$450), it would be completely unreasonable for you to give any significant credence to the proposition that a share of the total bill is \$450, despite the fact that this is what I, your peer, believe. In this case, it is the non-psychological considerations which swamp the psychological considerations into epistemic insignificance.

²⁶See Christensen (2007 pp.200-203) and Elga (forthcoming p.xx).

5. Considerations that Seem to Favor the Equal Weight View

5.1. Perceptual Judgments

As mentioned above, I believe that much of the appeal of The Equal Weight View derives from reflection on certain kinds of examples. In particular, The Equal Weight View can seem almost obviously or trivially correct when one reflects upon examples involving the conflicting perceptual judgments of individuals equally well-suited to make those judgments. Recall Case 1 from above: you and I, two equally attentive and well-sighted individuals, watch the horses cross the finish line from equally good vantage points. It looks to me as though Horse A finishes slightly ahead of Horse B, while it looks to you as though Horse B finishes slightly ahead of Horse A. The intuitive verdict: once we find that our initial judgments conflict, the uniquely reasonable course is for us to split the difference and retreat to a state of agnosticism about which of the two horses actually won the race.

I do not contest the intuitive verdict; indeed, I take it to be correct. What I do contest is the idea that the intuitive verdict has any tendency to support The Equal Weight View over The Total Evidence View. For when The Total Evidence View is correctly applied to Case 1, it too returns the intuitively correct verdict that you and I should abandon our original opinions and retreat to a state of agnosticism.

First, note that there are at least some cases in which The Total Evidence View will rationally require two individuals who begin with conflicting opinions to adopt a new opinion that is perfectly intermediate between their original opinions. Here is one such case:

Case 7. At time t_0 , you and I possess different evidence which bears on some hypothesis H . Your evidence suggests that H is true; my evidence suggests that it is false. Moreover, each of us responds to his evidence in a reasonable manner: you believe that H is true while I believe that it is false. At time t_1 , we encounter one another and pool our evidence. After doing so, our new total evidence does not favor H over not- H ; nor does it favor not- H over H .

Given that the total evidence available to us at time t_1 favors neither alternative over the other, an advocate of The Total Evidence View will maintain that we should suspend judgment. You should abandon your belief that the hypothesis is true while I should abandon my belief that it is false. In the light of our new total evidence, we should converge on the point that is intermediate between our original opinions. With respect to Case 7 then, The Total Evidence View will require us to respond in a way that is extensionally equivalent to the way that we would respond if we were both following a norm of ‘Split the Difference’.

Notice, however, that Case 7 is simply Case 1, abstractly described. As you and I watch the horses cross the finish line, it appears to me as though Horse A finishes just ahead of Horse B. To the extent that I have evidence for my judgment that Horse A finished ahead of Horse B, that evidence consists of my perceptual evidence: the fact that it *looks* or *appears* to me that Horse A finishes ahead, or that my visual experience represents Horse A as having finished ahead. In the absence of other evidence which bears on the question, it is at that point reasonable for me to believe that Horse A finished ahead of Horse B, since this is what my total evidence supports. Similarly, your initial judgment that Horse B finished just ahead of Horse A is a reasonable response to the evidence which you possess at time t_0 , viz. the fact that it looked or seemed to you as though Horse B finished just ahead of Horse A. At time t_1 , we compare notes: you learn that I think that Horse A won because that is how it looked to me; I learn that you think that Horse B won because that is how it looked to you. At this point, the total evidence that is available to each of us has changed in a rather dramatic way: I have gained evidence that suggests that Horse B won the race, while you have gained evidence that Horse A won the race. Moreover, given the relevant background assumptions and symmetries, it is natural to think that the total evidence that we now share favors neither the proposition that Horse A finished ahead of Horse B nor the proposition that Horse B finished ahead of Horse A. Thus, given our new total evidence, you and I should abandon our initial opinions about which horse won the race. The Total Evidence View, no less than The Equal Weight View, requires us to suspend judgment and retreat to a state of agnosticism in Case 1 and in cases of relevantly similar structure. Thus, it is a

mistake to think that such cases favor The Equal Weight View over The Total Evidence View.²⁷

5.2 The Analogy with Inanimate Measuring Devices

Recall Case 2 from above: you and I arrive at different views about the temperature by consulting our hitherto equally reliable thermometers; subsequently, we discover that our thermometers disagree. The intuitive verdict: we should abandon our original opinions about the temperature. In particular, it would be patently unreasonable for me to retain my original belief simply because that was what *my* thermometer indicated about the temperature. Indeed, I should give no more credence to what my thermometer says about the temperature than to what yours says, and *vice versa*. But what holds for the conflicting readings of equally reliable thermometers holds also for the conflicting judgments of individuals who are peers.

Response: I embrace the intuitive verdict about Case 2 but deny that this has any tendency to support The Equal Weight View over The Total Evidence View. Of course, the mere fact that a given thermometer is *mine* is no reason for me to think that it is more trustworthy than your thermometer, or for me to favor its indications over the indications of yours. But similarly, it is no part of The Total Evidence View that it is permissible for me to favor my original opinion over yours simply because the opinion in question belongs to me. On The Total Evidence View, once I discover that our original opinions

²⁷In general, it is important to distinguish between (i) cases in which multiple individuals have equally strong but different bodies of evidence and (ii) cases in which multiple individuals have equally strong bodies of evidence in virtue of sharing the same evidence. Splitting the difference will often be the reasonable response in the former kind of case, but this in itself has no tendency to show that the same is true in cases of the latter kind. Of course, a commitment to certain views about the nature of evidence might make it difficult, if not impossible, to consistently observe the distinction between (i) and (ii). For example, on a view of evidence according to which one's evidence ultimately consists of one's own private mental states, one never literally shares one's evidence with a peer; at best, one's evidence is similar in various salient respects to the evidence that one's peer possesses. Because this is the closest surrogate for genuinely sharing evidence in the literal sense, it becomes easy to conflate (i) and (ii). But such conflation should be resisted.

differ, it might very well be reasonable for me to adopt an opinion that is closer to my original opinion than to yours. But if so, that is because the opinion in question is best supported by what is now our total evidence. (By the same token, it might also be reasonable for me to adopt a new opinion that is closer to your original opinion than to mine.) Again, in such cases, whether the opinion that it is ultimately reasonable for us to hold is closer to your original opinion or to mine will typically depend on which one of us (if either) did a better job of responding to the first order evidence in arriving at his or her original opinion. When it is construed in this way, the thermometer analogy at best tells against the view which Elga refers to as ‘The Extra Weight View’.

The objector might insist that the thermometer analogy tells against The Total Evidence View as well. For consider: even if my thermometer is in fact functioning perfectly on a particular occasion and your thermometer is malfunctioning, it would still be unreasonable for me to favor what my thermometer says as long as I am in no position to appreciate these facts. Similarly, even if I have in fact evaluated the evidence correctly in arriving at a given belief and you have not, it would be unreasonable for me to favor my view over yours if I am in no position to appreciate that this is so.

Response: first, note that when the thermometer analogy is construed in this way, there are significant disanalogies between it and a case of peer disagreement. In the thermometer case, the two devices function as black boxes which simply output a given number; what takes place ‘underneath the hood’ (so to speak) is completely opaque to us. In terms of interpersonal disagreement, the closest analogue to this would be something like the following. I regard the two of us as generally reliable about some domain. I then discover that you hold a different view than I do about some issue in that domain, but I have no idea about how you arrived at your view or on what basis you currently hold it: for all I know, you might hold the view on the basis of relevant considerations of which I am unaware, considerations which, if presented to me, would result in a change in my view. On the other hand, it is also possible, for all I know, that you are unaware of relevant considerations which I possess (considerations which, if presented to you, would result in a change in your view). To the extent that I am ignorant of how you arrived at your view, or why you currently hold it, it is, I think, reasonable for me to treat our views even-handedly until learning more. (And of course, a proponent of The Total Evidence

View need not say otherwise.) In contrast, in cases of peer disagreement, the peers have access to the same body of evidence and are aware that this is so; as we have sometimes put it, they have ‘compared notes’. Thus, suppose that I confidently hold a philosophical thesis on the basis of arguments and considerations that are inadequate to support it.²⁸ I cite these arguments and considerations in defense of my thesis, attributing my confidence in the thesis to them. Recognizing my reasons as inadequate, you come away unimpressed. In these circumstances, there is a sense in which you have witnessed the malfunction occur. In terms of the thermometer analogy, it is as though one has the opportunity to open up the other person’s thermometer (as well as one’s own) and inspect how the temperature reading was arrived at. In cases of peer disagreement, one gets to go underneath the hood, as it were.²⁹

Still, one might think that this does not yet get to the heart of the matter. Even if the two peers are both fully aware of why the other believes as he or she does, and one of the two has in fact done a better job of evaluating their shared evidence, who is to say which of the two has done a better job? The peers themselves, one might think, are not in a position to justifiably make such judgments. Perhaps then the situation of the peers is analogous to the following: although one gets to open up the two thermometers and inspect how their conflicting readings are generated, one lacks the wherewithal to reliably discriminate a malfunctioning thermometer from one that is functioning properly. Similarly, a proponent of The Equal Weight View might claim that, even in a case in which one has in fact done a better job of evaluating the relevant considerations than one’s peer, one has no way of discriminating such a case from a case in which things are the other way around. Even if things are asymmetrical at the bottom level (one’s belief reflects the evidence better than one’s peer’s belief; one’s thermometer is in fact functioning better than the other person’s thermometer), things are symmetrical one level up: one has no justification for thinking that one’s belief better reflects the evidence, or

²⁸ Readers of the present paper still unconvinced of the truth of The Total Evidence View will no doubt find *this* thought experiment an eminently manageable one.

²⁹ For suggesting a response along these lines, I am grateful to Roy Sorensen (who originally put the thermometer objection to me, in conversation). I am unsure to what extent he takes the response to adequately defuse the objection.

that one's thermometer is the thermometer that is working properly. One is thus no more justified in thinking that one's own belief accurately reflects the evidence than one's peer is in thinking that *his* belief accurately reflects the evidence. Therefore, given the higher level normative symmetry, it would be unreasonable to favor one's own belief over the belief of one's peer. Favoring one's own belief would be reasonable only if one had some independent evidence that one's belief is more likely than one's peer's belief to be an accurate reflection of the evidence: for example, knowledge that one has outperformed the other person in relevant ways in the past. But *ex hypothesi*, no such independent evidence is available in a case in which the disagreement is a disagreement between peers. Thus—so the argument runs—even if one is in fact the person who has better evaluated the evidence on this particular occasion, one has no justification for thinking that this is so. Hence, one should split the difference.

However, the proponent of The Equal Weight View is not entitled to simply assume that things *are* symmetrical between us at the higher level. That is, the proponent of The Equal Weight View is not entitled to simply assume that you and I are equally well-justified in thinking that we have correctly responded to the evidence in a case in which you have done so and I have failed to do so. Of course, given that we are peers, neither of us possesses independent evidence which suggests that he is the one who has responded correctly. However, even in the absence of *independent* evidence, there is another possibility: namely, that when one correctly responds to a body of evidence, one is typically better justified in thinking that one has responded correctly than one is when one responds incorrectly. It is this thought which I will now explore and defend.

First, let us observe the unobvious point that, when one responds correctly to a body of evidence, one typically has some justification for thinking that one has responded correctly. In paradigmatic cases in which one takes up the view that is best supported by one's evidence, it is no mere accident that that one has done so (although lucky accidents are of course possible, they are atypical). Rather, one takes up the belief in question precisely because it is supported by one's evidence. Indeed, in a given case, one might very well take up the belief because one *recognizes* that this is what one's evidence supports. Plausibly, *recognizing that p* entails *knowing that p*. Assuming that that is so, then any case in which one recognizes that one's evidence supports a given belief is a

case in which one knows that one's evidence supports that belief. Clearly, if one knows that one's evidence supports a given belief, then one is *justified* in thinking that one's evidence supports that belief; if one were not justified, one would not know. But even if recognizing that *p* does *not* entail knowing that *p*, one would in any case not be able to recognize that *p* if one were unjustified in thinking that *p*. It follows immediately from this that whenever one recognizes that one's evidence supports such-and-such a conclusion, one is justified in thinking that one's evidence supports that conclusion.

Thus, in any case in which you hold a given belief because you recognize that this is what your evidence supports, not only is your belief a reasonable one, but you are also justified in believing an epistemic proposition to the effect that it is reasonable for you to hold that belief. Indeed, given that *recognizing that p* entails *being justified in believing that p*, you will be justified in believing that the evidence supports your view for as long as you continue to recognize that it does. Consider then a case in which you take up a belief in virtue of recognizing that that belief is what the evidence supports; I am not yet on the scene. Subsequently, you discover that I believe otherwise despite having been exposed to the same evidence. Presumably, the proponent of The Equal Weight View will claim that, once I arrive on the scene with my conflicting belief, you are no longer justified in believing that the original evidence supports the relevant proposition because you are now no longer in a position to recognize that it does: in effect, encountering a peer who thinks otherwise strips one of one's prior ability to recognize the bearing of the first-order evidence for what it is. This suggestion is perhaps not wholly implausible: the proponent of The Equal Weight View envisions the situation as one in which one loses knowledge that one previously possessed in virtue of acquiring misleading evidence. In general, this is perfectly possible. However, in order to preserve the higher level normative symmetry which lends The Equal Weight View its plausibility, it seems that the proponent of The Equal Weight View will have to maintain the following: when we met, the justification which I possessed for thinking that my original response to the evidence is reasonable was just as strong as the justification which you possessed for thinking that your original response to the evidence is reasonable. For if you were better justified in thinking that your response was reasonable than I was in thinking that my response was reasonable, then this would break the putative higher level symmetry and

provide a basis for favoring your original belief over mine. (Compare a situation in which you are better justified in thinking that your thermometer is functioning properly than I am in thinking that my thermometer is functioning properly.) Of course, because you originally recognized that the evidence supported your belief, there is, on the assumption that The Uniqueness Thesis is true, no possibility that I similarly recognized that the evidence supported my belief. At best, I mistakenly took the evidence to support my belief when it did not. In order to preserve normative symmetry at the higher level then, the proponent of The Equal Weight View will insist that

When you correctly recognize that the evidence supports p , you are no more justified in thinking that the evidence supports p than I am in thinking that the evidence supports not- p when I mistakenly take the evidence to support not- p .

Presumably, the reference to two people is inessential here. So the proponent of The Equal Weight View will also endorse

When you correctly recognize that the evidence supports p , you are no more justified in thinking that this is what the evidence supports than you would have been had you mistakenly taken the evidence to support not- p instead.

But these assumptions are quite dubious. In any case, they are not ones that a proponent of The Total Evidence View need or should accept. On its most formidable construal then, the argument from the analogy with inanimate measuring devices depends on assumptions that are at best controversial and for which no argument has been provided.

5.3 Downward Epistemic Push

Although the most recent objection is ultimately not compelling, it proceeds from a genuine insight which is worth making fully explicit. The insight in question is one that we have already briefly touched on above. It might be put like this: in general, what it is reasonable to believe about the world on the basis of one's evidence is constrained by

what it is reasonable for one to believe *about* one's evidence.³⁰ (Put otherwise: what it is reasonable for one to believe about the world is not independent of what it is reasonable for one to believe about what it is reasonable for one to believe about the world.)

Because this insight might naturally be taken to support The Equal Weight View, let us explore it a bit further here. Once again, let E represent one's total first order evidence with respect to H. Consider then the epistemic proposition that

E is good evidence for H

In section 4, it was argued that a proponent of The Total Evidence View should agree that higher order evidence about the quality of one's first order evidence typically makes some difference to what one should believe on the basis of that first order evidence. Inasmuch as this is so, she will agree with the proponent of The Equal Weight View that any evidence which bears on this epistemic proposition is also evidence which bears on H itself. In general, how confident one should be that H is true is tied to how confident one should be that the corresponding epistemic proposition is true. When one acquires reasons to increase one's confidence in the epistemic proposition, one acquires reasons to increase one's confidence in H; on the other hand, when one acquires reasons to decrease one's confidence in the epistemic proposition, one acquires reasons to decrease one's confidence in H. That this relationship exists might seem to strongly favor The Equal Weight View, or more generally, any view on which higher order evidence trumps lower order evidence. For even in a case in which E genuinely supports H, one will not be justified in believing H on the basis of E if one is not justified in believing the epistemic proposition that *E is good evidence for H*. And one's justification for believing this epistemic proposition would seem to be hostage to what one's peers think.

I think that this phenomenon of *downward epistemic push* is a genuine one. (If it were not, The Asymmetrical No Independent Weight View would be compelling.) However, a proponent of The Total Evidence View will insist that another point deserves equal

³⁰This point is well emphasized by Feldman (2006) and Christensen (forthcoming). Here is a representative quotation from the latter: 'the rationality of first order beliefs cannot in general be divorced from the rationality of certain second order beliefs that bear on the epistemic status of the first order beliefs' (p.xx).

emphasis, especially as it is apt to be given short shrift by one sympathetic to The Equal Weight View: there is also the opposite phenomenon, that of *upward* epistemic push. That is, a proponent of The Total Evidence View will insist (for the reasons given in 5.2 above) that when E is genuinely good evidence for H, this very fact will tend to justify belief in the epistemic proposition that *E is good evidence for H* for those with the relevant competence. It is not only that one's higher order evidence typically makes a difference to what one is justified in believing about the world; it is also the case that one's first order evidence makes a difference to what one is justified in believing about higher level epistemic matters.

Again, the phenomenon of upward epistemic push will be most visible in single person cases, where distracting complications are at a minimum. As argued above, any case in which one takes up a belief upon recognizing that that belief is supported by one's evidence is *ipso facto* a case in which one is justified in believing a corresponding epistemic proposition to the effect that *one's evidence supports that belief*. Moreover, it is implausible that every case in which one recognizes that a given belief is supported by one's first order evidence is a case in which one's recognition depends on one's having some independent, higher order evidence to the effect that one's evidence supports that belief. Rather, in some cases, one's recognition that one's evidence supports a given belief is based on an unmediated appreciation of that evidence itself. Thus, in such cases, one's first order evidence not only confirms the belief in question; it also confirms a proposition to the effect that it is reasonable for one to hold that belief.

I take the dialectical upshot of this picture to be the following: the proponent of The Total Evidence View can agree with the proponent of The Equal Weight View that facts about what one is justified in believing about the world are constrained by higher level facts about what one is justified in believing about one's evidence, while denying that this favors The Equal Weight View. (Indeed, the proponent of The Total Evidence will insist that when these matters are understood correctly, the picture that emerges will positively favor her own view.) In a case of disagreement in which one of two peers evaluates the first order evidence correctly, the proponent of The Equal Weight View will see a higher level symmetry and appeal to the link between levels in order to argue that neither peer is justified in favoring his original view once they compare notes.

(Symmetry at the higher level creates a symmetry at the lower level that otherwise would not have existed.) In contrast, a proponent of The Total Evidence View will contend that, in such a case, the peer whose view more accurately reflects the evidence will typically be better justified in thinking that his view is the one that is favored by the first order evidence. (Asymmetry at the lower level tends to create an asymmetry at the higher level, an asymmetry which otherwise would not have existed.)

No doubt, some will be extremely suspicious of the idea that the peer who gets the evidence right is typically better justified in thinking that he has done so than the peer who gets the evidence wrong. In fact, the attempt to exploit such suspicions is central to an argument for The Equal Weight View that has recently been offered by Adam Elga. The final argument for The Equal Weight View that I will consider then, is his.

5.4. A (No) Bootstrapping Argument for The Equal Weight View?

Elga argues as follows:

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 (pp.xx-xx).

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 ‘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 matters which of the parties has in fact done a better job evaluating the evidence.³¹ However, I believe that while Elga’s bootstrapping argument has considerable force against The Extra Weight View, it has little to none against The Total Evidence View.

In order to see this, let us 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

Crucially, however, this supposition is *not* something to which the 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 is not true, in general, that you should be more confident that your friend has made the mistake whenever the two of you disagree. In *some* cases, it might be reasonable for you to be more confident that your friend is the one who has

³¹Elga 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.xx).

made the mistake. But in other cases, it might be reasonable, given the total evidence available to you, to be more confident that *you* are the one who has made the mistake. On The Total Evidence View, it is not true that there is some general answer to the question of how confident you should be that it is 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 is reasonable to be that your friend has made a mistake is not something which 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 his view is false.

Consider another view rejected by Elga, The Extra Weight View. As interpreted by Elga, The Extra Weight View would license you in being extremely confident that you are a better evaluator than your friend simply by noting the many cases in which the two of you disagree. In a parallel manner, The Extra Weight View would license your friend in being extremely confident that *he* is the better evaluator by appeal to the very same disagreements. This seems odd (to say the least): the very same events are legitimately treated by you as confirming evidence for the claim that you are a better evaluator than your friend and by your friend as confirming evidence that he is a better evaluator than you. Moreover, even if you are in fact the inferior evaluator, and you consistently do a worse job evaluating the evidence on particular occasions, it will nevertheless be reasonable for you to conclude that you are superior to your friend *on the basis of those very cases*. (That is, it will be reasonable for you to conclude that you are a better evaluator of evidence on the basis of disagreements whose existence is underwritten by the fact that you have done a *worse* job than your friend has with respect to evaluating the evidence.) Here I agree with Elga: such a view makes it absurdly easy to arrive at evidence that one is a better evaluator. However, no similar absurdity follows from The Total Evidence View. It is true that the proponent of The Total Evidence View is committed to the following possibility: over time, you reasonably become quite confident that someone who you initially regarded as your peer is not your peer, on the basis of a large number of cases in which the two of you disagree. Consider, for example

Case 8. At the outset you regard your friend as your peer. Subsequently, however, many disagreements emerge. With respect to the vast majority of these disagreements, the position which you hold is in fact better supported by the available evidence than the position held by your friend. In these cases, your conviction that your friend's position is not adequately supported by his evidence is based on your own appreciation of that evidence, an appreciation which is more accurate than his. Over time, you thus become increasingly confident that you are a better evaluator of the evidence than your friend. You thus cease to regard your friend as your peer and conclude that your initial judgment to that effect was mistaken.

As Elga would have it, the proponent of The Total Evidence View is indeed committed to the possibility that such a change in view is reasonable in the envisaged circumstances. However, there is no absurdity here.

Elga's bootstrapping argument purports to establish that any view other than The Equal Weight View makes it too easy to reasonably conclude that you are a better evaluator than your friend. The danger in question is a real one: some views (e.g., The Extra Weight View) do fall victim to it. However, there is also the opposite danger: that a given view will make it *too difficult* to reasonably conclude that another person is not, contrary to what one initially thought, one's peer. Indeed, the line of argument offered by Elga seems to suggest something like the following: once you come to regard your friend as a peer about a given set of questions, it is not reasonable for you to demote him from the ranks of those to whom you accord that status on the basis of subsequent disagreements about those questions (rather, one would need to have independent evidence that you are a better evaluator than he is, evidence that is independent of the disputed issues themselves). But that seems too strong: to the extent that the argument purports to show this, the argument proves too much. For in some cases, it might very well be rational for you to conclude that your friend is not your peer after all, where your only basis for so concluding is the lack of judgment that he displays in subsequent cases in which the two of you disagree. The possibility of rationally downgrading someone from the status of peer in this way will be especially apparent in cases in which one's initial judgment that the other person is a peer was itself based on relatively insubstantial evidence. Consider for example

Case 9. At the first meeting of our seminar, I strike you as a perfectly reasonable and sensible person. For the most part, we find the same arguments and considerations persuasive. Even on those few occasions when we express different views, my view seems to you to be well within the bounds of reasonable opinion, no less than your own (suppose here that you do *not* accept The Uniqueness Thesis). On the basis of this first meeting then, you form the opinion that I am your peer.

In subsequent meetings of the seminar, however, you and I disagree often. Moreover, when we disagree, my views often seem to you to be based on relatively flimsy arguments; when I attempt to parry objections, what I say strikes you as weak and unresponsive, and so on. (Needless to say, I would dispute such assessments.) By the end of the semester, you no longer regard me as your peer.³²

Here, your revised estimate of my competence is based on your negative assessment of my performance in judging issues that are disputed between us. Moreover, the disputed issues are the very sorts of questions with respect to which you once reasonably took me to be a peer. Does this guarantee that it is unreasonable for you to demote me from the ranks of those to whom you accord such status? There is no such guarantee. On the other hand, there is also no guarantee that your demoting me *is* reasonable in the circumstances, given only the description of Case 9 offered above. Whether your demoting me is reasonable will typically depend on such things as whether my best attempts to parry objections are weak and unresponsive as you take them to be, or whether your conviction that they are weak and unresponsive is due to (e.g.) your being so dogmatically committed to the opposite conclusions that you fail to appreciate the merits of what I say. The more the former is the case, the more reasonable it will be for you to revise your estimate of my competence in a downward direction; the more the latter is the case, the less reasonable such revision is. Of course, from your perspective, it might be very difficult to tell which of these is the case. From the inside, a case in which you fail to appreciate the genuine merits of what I say on behalf of my view because of dogmatic commitment on your part might seem just like a case in which my defense is indeed without merit. But the fact that it might be difficult to tell which of these is the case does not mean that it makes no difference whether your revised estimate of my competence is based on your having recognized genuine shortcomings on my part, or whether it is instead an artifact of your own shortcomings. Here as elsewhere, there is no

³² This case was inspired by a similar example devised by Daniel Greco.

escape from the fact that one's judgment is fallible and subject to corruption in ways that tend to elude detection.

According to Elga, (i) the relevant kind of bootstrapping is never rationally permissible, (ii) The Equal Weight View proscribes such bootstrapping, and (iii) no other plausible view does so. He thus concludes that The Equal Weight View is true. I hold that, on the contrary, because there are at least some possible cases in which such bootstrapping clearly *is* permissible, no view which generally proscribes it can be correct. Hence, on the assumption that Elga is correct in thinking that The Equal Weight View generally proscribes such bootstrapping, we have arrived at another good reason for thinking that it is false.³³

³³A final word about Elga's blanket prohibition on the kind of bootstrapping at issue here. I suspect that this prohibition ultimately sits in at least some tension with moves that he makes elsewhere in the same paper. Consider, for example, his argument that The Equal Weight View does not require one to suspend judgment about all controversial issues (a conclusion which would be, he thinks, an absurd consequence). In attempting to block this 'problem of spinelessness', Elga emphasizes that we should not overestimate how often we find ourselves in disagreements with those whom we take to be our peers, inasmuch as I will not consider you my peer with respect to a question if you disagree with me about too many surrounding issues. (Here, it seems, Elga would permit one to appeal to one's own beliefs in order to conclude that someone who disagrees with sufficiently many of those beliefs is not a peer.)

However, there is at least a certain awkwardness in attempting to combine (i) a blanket prohibition against bootstrapping one's way to the conclusion that someone who one initially took to be a peer is not a peer, and (ii) permitting one to appeal to one's own beliefs in order to deny the status of peer to another person so long as one has not yet accorded him that status. For suppose that you and your friend disagree, not only about the moral permissibility of abortion, but also about many surrounding issues, e.g., whether human beings have souls, whether it is permissible to withhold treatment from certain terminally ill infants, whether rights figure prominently in a correct ethical theory, and so on. (I take this example directly from Elga; see pp.xx-xx.) According to Elga, because your friend has by your lights come to the wrong conclusion about this entire cluster of closely related issues, you (reasonably) do not consider him your peer, and thus, you are not required to split the difference with him about these issues. However, even this amount of disagreement is presumably compatible with a large amount of agreement concerning moral matters. Bearing this mind, consider two cases. In Case A, you first discover all of the issues with respect to which you and your friend disagree; you thus conclude, reasonably, that your friend is not your peer with respect to difficult moral questions. You can therefore retain your original views, since you are not rationally required to split the difference and retreat to a state of agnosticism. Later, it

6. The Total Evidence View: Concluding Remarks

In the course of laying out the core ideas of The Total Evidence View, we have repeatedly compared and contrasted it with various alternatives, especially The Equal Weight View. As we have seen, The Total Evidence View and The Equal Weight View yield the same verdict when applied to some cases of disagreement; in others, the two views differ sharply in what they require of the disputing parties. In still other cases, I think that it is somewhat unclear how far the two views differ (if they differ at all), inasmuch as it is somewhat unclear how The Total Evidence View should be applied.

Consider, for example, cases in which individuals of apparently similar mathematical abilities arrive at different conclusions on the basis of temporally extended token processes of calculation.³⁴ In some cases of this sort, I believe that The Total Evidence View, when properly interpreted, will yield the verdict that the individuals should split the difference. Consider for example

Case 10. You and I add a series of ten three digit numbers in our heads. A third party calls out the numbers, one after the other. Each of us keeps a running tally, adding the numbers as we go, not attempting to keep track of any particular number in the sequence after it has been added to the running total. We know that, when we have

emerges that you and your friend also agree on many moral issues; you thus view these cases as ones in which someone who is not your peer nevertheless manages to arrive at the correct conclusions. In contrast, in Case B you first happen to discover all of the moral issues with respect to which you and your friend agree; you thus (reasonably) conclude that your friend is your peer with respect to difficult moral questions. Later, all of the disagreements emerge. However, because you have already granted your friend the status of peer, it is no longer permissible, given the prohibition on bootstrapping, to appeal to these disagreements as a legitimate basis for demoting him or for discounting his opinion: at this point, you are already on the normative hook, as it were. Ultimately, you are fully aware of all of your friend's opinions in both cases; the only difference is that in Case B, you are rationally required to suspend judgment on all of the disputed questions while in Case A, you are not. But intuitively, that seems wrong. For the only underlying difference between the two cases is the *order* in which you learn of your friend's opinions.

³⁴I take this to be among the issues raised by Christensen's (2007) 'Restaurant Case'.

played this game in the past, we have made a more or less equal number of mistakes. This time, I arrive at the number 5,863 and you arrive at the number 5,883. Once we discover that we have arrived at different answers, how should we respond?

One might think that, on The Total Evidence View, although both of us should be less confident of our original answers, whichever one of us has in fact performed the calculation correctly is rationally entitled to be more confident of her answer than the other person. However, in this case I believe that there is a stronger, countervailing pressure from within The Total Evidence View which militates in favor of splitting the difference. On The Total Evidence View, what it is reasonable for us to believe always depends on the total evidence which we possess. Once we learn that we have arrived at different answers, then, given that we have not retained specific information about the original numbers (and so are not in a position to reconstruct our original reasoning), it seems as though the total relevant evidence available to us consists of the fact that (i) I arrived at the 5,863 and (ii) you (a person of apparently similar mathematical abilities) arrived at the number 5,883. Plausibly, this evidence does not favor either answer over the other. (Compare the question of what it would be reasonable for a third party to believe, a person who was not present when the original numbers were called out and who knows only (i) and (ii)). Hence, it seems like The Total Evidence View requires us to split the difference in these circumstances.

Suppose that we subsequently go over our reasoning together, step by step. (The person who originally called out the numbers recorded them on a list.) If we are in fact both competent at arithmetic, the person who originally made the mistake will presumably correct his or her error and adopt the true view. Imagine, however, that our dispute turns out to have the kind of persistent, intractable character that many philosophical disputes seem to possess: despite explicitly rehearsing our chains of reasoning in a public manner, both of us continue to think that his or her original answer is correct. Of course, because of the quasi-algorithmic character of arithmetical reasoning, it is difficult to imagine that a dispute of this sort could persist. But if such were to happen, I think that it is reasonable for the person whose answer is based on the sound mathematical reasoning to be relatively confident of her answer. She occupies a superior epistemic position compared to the person whose answer is based on the

unsound reasoning. Yet it is far from clear how we should think about the mathematical evidence in such a case, or in cases of calculation more generally. To the extent that the relevant notion of mathematical evidence remains unclear, there will be many cases of calculation in which it is unclear how exactly The Total Evidence View applies.

Still, even if there are unclarities about how The Total Evidence View applies in particular cases, the view is at least clear enough to be controverted. I close by commenting briefly on two features that are especially likely to arouse suspicion in certain quarters.

First, a central feature of the view is that the reasonableness of the parties in a case of peer disagreement will typically depend on whose opinion better reflects the first order considerations relevant to their dispute. When one is responding correctly to the evidence, one is typically in a stronger position vis-à-vis those who think otherwise than when one merely takes oneself to be responding correctly to the evidence. Of course, there is no magic red light that illuminates when one responds to the evidence correctly, no warning bell which sounds when one does not. Indeed, as a phenomenological matter, there might be no introspectible difference between how things seem when one is responding correctly and how things seem when one is not. Given this, how can a decision to adopt The Total Evidence View (or any view which shares this central feature) reflect anything other than a meta-epistemological commitment to externalism about justification, with The Equal Weight View left as the view of choice for those with more internalist sympathies?

However, the classification of The Total Evidence View as an ‘externalist’ as opposed to an ‘internalist’ view is not a happy one. On the Total Evidence View, what it is reasonable for one to believe always depends on one’s total evidence, and only considerations of which one is aware are eligible for inclusion in one’s total evidence. (Relevant considerations that are known to others but of which one is unaware make no difference to what it is reasonable for one to believe.) In this crucial respect, The Total Evidence View resembles epistemological views that are paradigmatically ‘internalist’. Of course, in a case in which one’s view is not adequately supported by one’s evidence, there might be nothing which indicates that this is so (that is, nothing *else*, beyond the evidence itself, which *ex hypothesi* one has misjudged). Because of this, when one’s

judgment as to the epistemic status of some belief that one holds is faulty, there is nothing which guarantees that this fact will be revealed by further reflection, no matter how conscientiously such reflection is conducted. But it is dubious that this last feature is avoided by any plausible view about justification, including paradigmatically internalist ones.

A related but ultimately deeper source of resistance to The Total Evidence View is the tendency to identify *good evidence* with *potentially persuasive evidence*. Relevant here is the seductive appeal of what Timothy Williamson has dubbed ‘the dialectical conception of evidence’ (Williamson 2004 and forthcoming). Recall from above the way in which a proponent of The Total Evidence View will differ from Elga regarding the legitimacy of a certain kind of bootstrapping. On The Total Evidence View, even if you reasonably take me to be your peer with respect to a given class of questions, you might later reasonably revise this judgment and conclude that I am not your peer on the basis of how I answer those very questions.³⁵ Consider a case in which you demote me in this way, and imagine that I subsequently learn that you no longer consider me a peer. Naturally enough, I inquire about your basis for demoting me. What evidence have you gained, since the time when you reasonably took me to be your peer, which suggests that I am not? In response to this query, you might cite one of those issues with respect to which I have by your lights misjudged the evidence. That is, for some issue about which we disagree, you might say the following:

With respect to this particular issue, the view which I hold is adequately supported by the evidence while the view which you hold is not. Hence, this case is a piece of confirming evidence for the claim that I am a better evaluator of evidence than you are.

Needless to say, I will not be impressed with this response, or agree that you have succeeded in providing any evidence of your superiority. By my lights, this is mere assertion on your part (and indeed, false assertion at that). Moreover, it should come as no surprise to you that this will be my reaction. Given that the question of what the

³⁵ Again, whether it is reasonable for you to downgrade me in this way will typically depend on whether you are correct in your assessment that my performance in evaluating our shared evidence has been inferior to your own.

evidence supports is a matter that is contested between us, the question of which one of us has done a better job of evaluating that evidence will also be a contested matter. For this reason, it would be completely unreasonable on your part to expect me to treat what you say here as evidence of your superiority. Indeed, it would be pointless for you to offer these alleged facts to me as evidence of your superiority, for from my perspective they are not facts at all. But if it would be pointless for you to cite these alleged facts as evidence in response to my request, then it would be *dialectically inappropriate* for you to do so. According to the dialectical conception of evidence, only considerations that it would be dialectically appropriate to cite as evidence are genuine evidence. On the dialectical conception of evidence then, the alleged facts which you cite as evidence of your superiority do not constitute such evidence, and thus have no tendency to justify your newfound belief in your superiority (even if what you say is true, and the alleged facts to which you appeal are genuine facts).

Moreover, the question of whether genuine evidence must consist of considerations that it would be dialectically appropriate to cite as such is of quite general relevance to the topic of peer disagreement. *Ex hypothesi*, when you and I disagree about an issue with respect to which we are peers, the fact that you do not share my view is not due to my having access to some crucial piece of evidence to which you lack access. One might think that this already renders problematic the claim that my evidence suffices to justify my belief: if my evidence were really sufficient to justify my believing as I do, wouldn't that evidence be enough to persuade you as well? Given that the evidence on which I base my belief does not persuade you, can't we conclude from this that my evidence is not sufficient to justify my belief after all? But in any case, when faced with someone who shares my evidence yet remains unconvinced, it would be pointless for me to simply recite the same considerations again and claim that that is why my view is the reasonable one to hold. Inasmuch as it would be pointless for me to cite these considerations in this way, it would be dialectically inappropriate for me to do so. Thus, if in order for it to be the case that my evidence genuinely favors my view over yours, it must be the case that it would be dialectically appropriate for me to cite my evidence as favoring my view over yours, then my evidence does not favor my view over yours.

I believe that we have good reasons to reject the dialectical conception of evidence. One might have good evidence that some claim is true even if one has no potentially persuasive evidence, or evidence that it would be dialectically appropriate to cite as such.³⁶ Indeed, I believe that one can have good evidence that some claim is true even if one has no evidence that it would be dialectically appropriate to offer to a person who is in general *no less reasonable than oneself*. Because lapses and blindspots are possible, the fact that a generally reasonable person fails, even repeatedly fails, to be persuaded of some conclusion by a body of evidence does not suffice to show that that evidence is inadequate to justify belief in that conclusion. Of course, once it is clear that the person does not find the considerations on offer persuasive, the dialectically appropriate course is to seek new considerations which might inspire conviction rather than to simply recite the original considerations yet again. But it does not follow from this that the original considerations were themselves inadequate to justify belief in the conclusion. The link between genuine evidence and potentially persuasive evidence is not as close as the dialectical conception of evidence suggests. At best, what is true is a relatively trivial claim: genuine evidence is evidence which will tend to persuade someone who will respond to that evidence in a fully reasonable manner. Still, it must be admitted that the dialectical conception of evidence is not wholly without its appeal. Inasmuch as this is so, a defense of The Total Evidence View more thorough than the one offered here will account for this appeal in a way that reveals it to be spurious.

³⁶ See especially Williamson (2004). As Williamson notes, acceptance of the dialectical conception of evidence would hand a cheap and sweeping victory to the crudest of skeptics. Thus, against a skeptic who consistently maintained that nothing is evidence for anything else, anything that one might offer as evidence would fail to qualify as such when judged by the dialectical standard; if meeting the dialectical standard was necessary for something to count as genuine evidence, one would have no genuine evidence at all when in the presence of such a skeptic. But surely this is incorrect. One can have genuine evidence, i.e., evidence which tends to justify one's beliefs, even when one has no evidence that it would be dialectically appropriate to offer.

On the dangers of not recognizing the distinction in question, see also James Pryor (2004).

References

- Christensen, David (2007). "Epistemology of Disagreement: the Good News", in *The Philosophical Review*, volume 116, no.2, pp.187-217.
- Christensen, David (forthcoming). "Does Murphy's Law Apply in Epistemology? Self-Doubt and Rational Ideals". To appear in *Oxford Studies in Epistemology 2* (Oxford University Press, Oxford).
- Edwards, Paul (1967). "Common Consent Arguments for the Existence of God". In Paul Edwards (ed.), *The Encyclopedia of Philosophy*, volume 2, pp.147-155.(Macmillan, NY).
- Elga, Adam (forthcoming). "Reflection and Disagreement". To appear in *Nous*.
- Feldman, Richard (2003). *Epistemology*. (Prentice Hall, Upper Saddle River, NJ).
- Feldman, Richard (2005). "Respecting the Evidence" in *Philosophical Perspectives, vol.19: Epistemology* (Blackwell, Oxford) pp.95-119.
- Feldman, Richard (2006). "Epistemological Puzzles About Disagreement" in Stephen Hetherington (edd.) *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*.
- Goldman, Alvin (2001). "Experts: Which Ones Should You Trust?" *Philosophy and Phenomenological Research*, Volume 63, 2001.
- Greco, Daniel (unpublished). "Disagreement and the Equal Weight View". Unpublished senior thesis, Princeton University, 2006.
- Kelly, Thomas (2005). "The Epistemic Significance of Disagreement", in Tamar Szabo Gendler and John Hawthorne (eds.), *Oxford Studies in Epistemology*, volume 1, (Oxford University Press, Oxford) pp.167-196.
- Kelly, Thomas (unpublished). "Comments on White, 'Epistemic Permissiveness'". Remarks delivered at the 2005 Eastern Division meetings of the American Philosophical Association. Available online at xxxx.
- Plantinga, Alvin (2000a). "Pluralism: A Defense of Religious Exclusivism", in Philip L. Quinn and Kevin Meeker (eds.), *The Philosophical Challenge of Religious Diversity*, (Oxford University Press, Oxford), pp.172-192.

Plantinga, Alvin (2000b). *Warranted Christian Belief* (Oxford University Press, Oxford).

Pryor, James (2004) "What's Wrong with Moore's Argument?" *Philosophical Issues* 14 (2004).

Rosen, Gideon (2001). "Nominalism, Naturalism, Epistemic Relativism", *Philosophical Perspectives* 15 (Blackwell, MA), pp. 69-91.

Sosa, Ernest (1999). "Skepticism and the Internal/External Divide", in John Greco and Ernest Sosa (eds.), *The Blackwell Guide to Epistemology* (Blackwell, MA).

Sosa, Ernest (2002). "Privileged Access", in Quentin Smith and A. Jokcik (eds.) *Consciousness: New Philosophical Essays* (Oxford University Press, Oxford), pp.273-295.

Sosa, Ernest (2007). *A Virtue Epistemology* (Oxford University Press, Oxford).

van Inwagen, Peter (1996). "It is Wrong, Everywhere, Always, and for Anyone, to Believe Anything up Insufficient Evidence", in Jeff Jordan and Daniel Howard-Snyder (eds.) *Faith, Freedom, and Rationality: Philosophy of Religion Today* (Rowman and Littlefield, London), pp.137-153.

Wedgwood, Ralph (forthcoming). *The Nature of Normativity*.

White, Roger (2005). "Epistemic Permissiveness" in *Philosophical Perspectives, vol.19: Epistemology* (Blackwell Publishers, Oxford), pp.445-459.

Williamson, Timothy (2000). *Knowledge and Its Limits*. (Oxford: Oxford University Press).

Williamson, Timothy (2004). "Philosophical 'Intuitions' and Skepticism About Judgment" *Dialectica* 58, 1, pp.109-153.

Williamson, Timothy (2007) *The Philosophy of Philosophy*. (Blackwell, Oxford).