

Peer Disagreement and Higher Order Evidence

Thomas Kelly

1.

My aim in this article 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 that 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.

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.¹ 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:

[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)²

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. For each of us, suspension of judgment is now the uniquely reasonable attitude. 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.

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 sixty-eight degrees, and so immediately take up the corresponding belief. Meanwhile, you consult your thermometer, find that it reads seventy-two 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?³

I take it as obvious that in these circumstances I should abandon my belief that it is sixty-eight degrees and you should abandon your belief that it is seventy-two 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 or her 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 sixty-eight or seventy-two 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 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*.⁴ 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 on 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 own 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, on 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 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.⁵ 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 on 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 equal weight view 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 behalf of the symmetrical no independent weight view.

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. It is sometimes defended in contexts in which the propositional attitude of belief is treated as an all-or-nothing matter: for any proposition one considers, one has in effect three doxastic options—one either believes the proposition, disbelieves the proposition, or suspends judgment as to its truth.⁶ 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 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 that treats belief as an all-or-nothing matter. Thus, consider a possible world that consists of two peers, one of whom is a theist and the other 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 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 that recognizes more fine-grained psychological states. Let us then adopt the standard Bayesian convention according to which the credence 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, the import of the equal weight view is clear: on 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," that is, cases in which the relevant proposition is initially either believed or disbelieved by the parties,⁷ it seems that an advocate of the equal weight view still has strong reasons to insist on a framework that 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, namely 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 that 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 it is plausible that 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.

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 I hold about H is within the range of perfectly reasonable opinion, as is the opinion you hold. Moreover, both of us have recognized this all along. Why then would we be rationally required to change?

Someone 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.⁸

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

How plausible is the uniqueness thesis? For my part, I find that its intuitive plausibility depends a great deal on how one thinks of the psychological states to which it is taken to apply. The uniqueness thesis seems most plausible when one thinks of belief in a maximally coarse-grained way, as an all-or-nothing matter.⁹ 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 that employs a fine-grained notion of belief.

Some philosophers find it pretheoretically obvious that the uniqueness thesis is false.¹⁰ Many others accept substantive epistemological views from which its falsity follows.¹¹ 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 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 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 that 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 most 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.¹² 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, 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 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. This is the reasonable level of confidence for both of us to have at time t_1 . 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 that 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.¹³ 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 on the total evidence we possess at that point. Let's call the total evidence we possess at time t_1 E^* . What does E^* include? Presumably 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 t_1 supervenes on how you and I respond to E at time t_0 . With respect to playing a role in determining what is reasonable for us to believe at time t_1 , E gets completely swamped by purely psychological facts about what you and I believe.

I find this consequence 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.

3.1. 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 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 that completely abstracts away from facts about which beliefs are

better supported by one's evidence, and that 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 that make such a view unattractive are shared by the equal weight view.

3.2. 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 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 t_1 , 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 t_1 is the attitude it is reasonable for us to take, despite the poor job 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.

Can the equal weight view be interpreted in such a way that it does not allow for such bootstrapping? A proponent might suggest the following: in response to peer disagreement, one is *rationally required* to split the difference, but it does not follow that the opinion at which one arrives by doing so is reasonable. Rather, splitting the difference is a *necessary but insufficient condition* for the reasonableness of the opinion at which one arrives. In order for that opinion to be reasonable, one must not only have arrived at it by splitting the difference, but one must have correctly responded to the original evidence as well. Thus, peers who scrupulously adhere to the equal weight view will wind up with reasonable opinions if they begin from reasonable opinions, but not if they begin from unreasonable opinions. In this way, the current bootstrapping objection is apparently blocked.

However, this proposed interpretation runs into serious problems elsewhere. Consider again Case 4, in which you but not I respond to the original evidence E in a reasonable manner. At time t_1 , we discover our disagreement and split the difference, converging on a credence of .5. On the present proposal, your credence of .5 is perfectly reasonable, since you have responded to the evidence correctly at every stage. On the other hand, my credence of .5 is *not* reasonable, since I misjudged the original evidence; the mere fact that I respond appropriately to your opinion by splitting the difference is not sufficient to render the opinion at which I thereby arrive reasonable. But here something seems to have gone wrong. After all: notice that at time t_1 , you and I have exactly the same evidence

that bears on H (viz. E, plus our knowledge of how each of us originally responded to that evidence), and we invest exactly the same credence in H on the basis of that evidence (viz. .5), yet your credence is reasonable on the evidence while mine is not. That seems wrong.¹⁵ Thus, although this interpretation of the equal weight view manages to avoid the charge of bootstrapping, it is untenable on other grounds. I therefore set it aside.

3.3. Even Easier, and More Implausible, Bootstrapping: Single Person Cases

On the equal weight view, the evidence that 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 nonpsychological evidence.¹⁶ 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, namely a case in which the psychological evidence is all the evidence 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 I 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 nonpsychological evidence has dwindled to nothing). At the other end of the spectrum are cases in which all of one's evidence is nonpsychological (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 nonpsychological 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 nonpsychological 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 nonpsychological 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 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 that bears on the truth of the hypothesis. At time t_1 , then, one's total evidence consists of one's original body of nonpsychological evidence E , plus a single piece of psychological evidence, namely 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 nonpsychological 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 .3 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 nonpsychological evidence E
 The fact that one believes the hypothesis to degree .7
 The fact that one's peer believes the hypothesis to degree .3

According to the equal weight view, one should split the difference with one's peer and believe the hypothesis to degree .5 at time t_2 . I have criticized the view on the grounds that it implausibly suggests that the psychological evidence swamps the nonpsychological evidence in these circumstances. At present, however, I want to inquire about what a proponent of the equal weight view should say about what one is rationally required 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 nonpsychological 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 one is rationally required to believe is entirely determined by 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 one is rationally required to believe will be entirely determined by 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 . . . and so on. In general, for any time t_n , a proponent of the equal weight view will hold that what one is rationally required to believe is entirely fixed by the opinions of the n peers. Why then should things be any different back at time t_1 , when the number of peers is 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 nonpsychological 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 nonpsychological 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 rationally required credence would be .6 or .8. On the picture of evidence suggested by the equal weight view, the distinction between believing and believing rationally seems to collapse in cases in which one is aware of what one believes but unaware of what others believe.

Here I 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 in which the qualified teacher of philosophy stands 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 overestimate 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 determined by giving equal weight to the opinions of the peers; 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.¹⁷ Bare truths about who has in fact manifested his or her 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 multiperson 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 nonpsychological evidence he or she possesses). Rather, on the suggested picture, if I am generally competent in the way 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 I do. But this contradicts our initial assumption, namely 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.

3.4. 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 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 percent 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 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 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 the judgments of some other mind come to play the role of evidence does not differ from the way 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 nonpsychological 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 nonpsychological 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 nonpsychological evidence, but where your psychological evidence is exhausted by the fact that (1) you yourself are confident that the proposition is true, and (2) some peer is equally confident that the proposition is false. Again, on the equal weight view, you should split the difference with your peer and retreat to a state of agnosticism; in effect, one ought to give no weight to the nonpsychological 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 nonpsychological evidence for nothing, *given that the way the two kinds of evidence qualify as such is exactly the same?*

4. THE TOTAL EVIDENCE VIEW

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.

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 if you have evaluated the evidence correctly and I have not. 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, namely, 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 his or her guns, and the other peer should convert, since the opinion in question is the one that 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 case you would have been in the same state of practical certainty on entering the office of the second colleague 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 on entering the office of the second colleague. And if it is rationally permissible to give zero weight to *his* opinion . . . and so on.

Moral: the fact that a peer believes differently can make it rationally incumbent on 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.¹⁸ 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 the 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.¹⁹

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 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.

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 one's practical reasoning. 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 on 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.²⁰

Where does this leave us?

In section 3, 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 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 that bear on this issue here. Consider again the kind of case I 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 that 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**:

- (1) Our original evidence *E*
- (2) The fact that you are quite confident that *H* is true
- (3) 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 ((2) and (3)) 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 that tends to confirm *H*, another that 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, namely those 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 (2) and (3) 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 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 that consists of higher order psychological evidence increases, and the proportion of the total evidence that 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 opinion is distributed among the peers. That is, it holds for cases in which peer opinion is evenly divided and 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 that 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 nonpsychological 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. Earlier, 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 section 3.2). 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 arrive at the same conclusion by independently committing the same fallacy, it will typically be reasonable for them to believe that conclusion on 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 that 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 toward 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 (2005a), I took a dim view of the suggestion that this fact would suffice to 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 nonskeptics. 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 that 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.²¹ 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.

(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.)²²

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 I 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 nonindependently. 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.

The same holds for cases in which there is widespread disagreement but the members of the contending factions have not arrived at their opinions independently. In an interesting recent essay, G. A. Cohen (2000) 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. (18, emphases in original)

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. 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.²³

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 (2007, pp. 490–91). 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 answer of \$450, an amount that 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 nonabsurd answer evenhandedly.

Unsurprisingly, both Christensen and Elga have interesting and detailed stories to tell about why, in these but not in otherwise similar cases, one need not give any weight to the view of one's peer.²⁴ I will not pause to evaluate the specifics of their respective proposals; here I 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 that bear on the question at issue (e.g., your knowledge that the total bill

is n , a number that 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 nonpsychological considerations that swamp the psychological considerations into epistemic insignificance.

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 on examples involving the conflicting perceptual judgments of individuals equally well suited to make those judgments. Recall Case 1: 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 began 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 that 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 or her 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 that 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 that you possess at time t_0 , namely 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, while you have gained evidence that Horse A won. 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.²⁵

5.2. 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. (2007, p. 487)

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.²⁶ 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, namely the situation in which you and your friend each evaluates 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 percent 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* percent 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 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 that floats entirely free of the evidence on which he bases his opinion. Thus, since the proponent of the total evidence view would not accept the supposition from which Elga derives the absurd consequence, the *reductio ad absurdum* on offer cannot show that her view is false.

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): exactly the 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 that 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 that 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.²⁷

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 is instead an artifact of your own shortcomings. Here as elsewhere, there is no escape from the fact that one's judgment is fallible and subject to corruption in ways that tend to elude detection.

According to Elga, (1) the relevant kind of bootstrapping is never rationally permissible, (2) the equal weight view proscribes such bootstrapping, and (3) 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 that 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.

Notes

This essay is something of a sequel to Kelly (2005a). 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 essay were presented at New York University, MIT, Rutgers University, Brown University, Princeton University, and the University of California at Irvine; I am grateful to the audiences present on those occasions. In addition, I would like to thank Aaron Bronfman, David Christensen, Adam Elga, Hartry Field, Allan Gibbard, Margaret Gilbert, Daniel Greco, Aaron James, Jim Joyce, Sarah McGrath, Philip Pettit, Jim Pryor, Walter Sinnott-Armstrong, Roy Sorensen, and Ernest Sosa for helpful conversations on the topic.

*Editorial note: this essay is a significantly abridged version of one by the same title published in *Disagreement*, edited by Richard Feldman and Ted Warfield Oxford: Oxford University Press (2010).

1 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 with reliability in accurately answering metaphysical questions.)

2 Compare Feldman (2003): after reviewing a number of examples of the kind at issue here, Feldman draws the conclusion: “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 (2007), 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 I will put forth differs from theirs, I have learned much from each of these authors.

3 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).

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

5 Compare “the Extra Weight View” discussed by Elga (2007), who argues against it.

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

7 Again, this is characteristic of Feldman’s work on the topic.

8 “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 it and offers some considerations for thinking that it is true. White (2005) argues for it at length but stops short of endorsing it.

9 Most 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.

10 Here, 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)

11 See, e.g., the brief survey in White (2005, pp. 445–46).

12 I 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 Kelly (2005b).

13 Is there some way of interpreting the equal weight view so that it does not have the consequence in question? On this possibility, see section 3.2.

14 The objection raised in this section is due to Aaron Bronfman. I utilize it here with his permission.

15 In any case, I take it that it is not an acceptable consequence for an evidentialist like Feldman, who explicitly maintains that what one is justified in believing

at any given time supervenes on what evidence one possesses at that time. See Conee and Feldman (2004), especially essay 4 and the introduction.

16 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 "nonpsychological evidence" as "doxastic evidence" and "nondoxastic evidence" in what follows.

17 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 that distinguishes it from the alternative view that I will offer and leaves it vulnerable to the current objection.

18 See the lucid and illuminating discussions of this point in Christensen (2007, 2008).

19 In Kelly (2005a), 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 at 185–90). 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 of the evidence on which the person bases her belief. (Foremost among my reasons for skepticism: the "double-counting" argument rehearsed at 187–88.) What I say here constitutes a departure from the earlier skeptical attitude: on this 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.

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

21 On the importance and nature of independence, see especially the illuminating discussion in Goldman (2001, pp. 150–56). In that essay Goldman is specifically concerned with the interesting question of how a nonexpert should respond to disagreement among the experts, but the analysis of independence he offers would seem to be highly relevant to a host of other important issues in social epistemology as well.

22 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).

23 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).

24 See Christensen (2007, pp. 200–203) and Elga (2007, p. 491).

25 In general, it is important to distinguish between (1) cases in which multiple individuals have equally strong but different bodies of evidence, and (2) 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 (1) and (2). 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 one's peer possesses. Because this is the closest surrogate for genuinely sharing evidence in the literal sense, it becomes easy to conflate (1) and (2). But such conflation should be resisted.

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

Again, this absurdity is independent of who has in fact evaluated the claims properly. Even if in fact you have done a much better job than your friend at evaluating the claims, simply comparing your verdicts to those of your friend gives you no evidence that this is so. (2007, p. 487)

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

References

- Christensen, David (2007). "Epistemology of Disagreement: The Good News." *Philosophical Review* 116, no. 2, 187–217.
- Christensen, David (2008). "Does Murphy's Law Apply in Epistemology? Self-doubt and Rational Ideals." In *Oxford Studies in Epistemology*, vol. 2. Oxford: Oxford University Press, pp. 3–31.
- Cohen, G. A. (2000). *If You're an Egalitarian, How Come You're So Rich?* (Cambridge, MA: Harvard University Press).
- Conee, Earl, and Feldman, Richard (2004). *Evidentialism*. Oxford: Oxford University Press.

- Edwards, Paul (1967). "Common Consent Arguments for the Existence of God." In Paul Edwards (ed.), *The Encyclopedia of Philosophy*, vol. 2 New York: Macmillan, pp. 147–55.
- Elga, Adam (2007). "Reflection and Disagreement." *Nous* 41, no. 3, 478.
- Feldman, Richard (2003). *Epistemology*. Upper Saddle River, N.J.: Prentice Hall.
- Feldman, Richard (2005). "Respecting the Evidence." In *Philosophical Perspectives*, vol. 19, *Epistemology*. Oxford: Blackwell, pp. 95–119.
- Feldman, Richard (2006). "Epistemological Puzzles about Disagreement." In Stephen Hetherington (ed.), *Epistemology Futures*. Oxford: Oxford University Press, pp. 216–36.
- Feldman, Richard (2007). "Reasonable Religious Disagreements." In Louise Antony (ed.), *Philosophers without Gods: Meditations on Atheism and the Secular Life* (Oxford: Oxford University Press), pp. 194–214.
- Goldman, Alvin (2001). "Experts: Which Ones Should You Trust?" *Philosophy and Phenomenological Research* 63: 85–110.
- Greco, Daniel (2006). "Disagreement and the Equal Weight View." Senior thesis, Princeton University.
- Kelly, Thomas (2005a). "The Epistemic Significance of Disagreement." In Tamar Szabo Gendler and John Hawthorne (eds.), *Oxford Studies in Epistemology*, vol. 1. Oxford: Oxford University Press, pp. 167–96.
- Kelly, Thomas (2005b). "Comments on White, 'Epistemic Permissiveness.'" Remarks delivered at the annual meeting of the Eastern Division of the American Philosophical Association. Available at www.princeton.edu/~tkelly/.
- Plantinga, Alvin (2000a). "Pluralism: A Defense of Religious Exclusivism." In Philip L. Quinn and Kevin Meeker (eds.), *The Philosophical Challenge of Religious Diversity*. Oxford: Oxford University Press, pp. 172–92.
- Plantinga, Alvin (2000b). *Warranted Christian Belief*. Oxford: Oxford University Press.
- Rosen, Gideon (2001). "Nominalism, Naturalism, Epistemic Relativism." *Philosophical Perspectives* 15, 69–91.
- Sosa, Ernest (1999). "Skepticism and the Internal/External Divide." In John Greco and Ernest Sosa (eds.), *The Blackwell Guide to Epistemology*. MA: Blackwell.
- Sosa, Ernest (2002). "Privileged Access." In Quentin Smith and A. Jokcik (eds.), *Consciousness: New Philosophical Essays*. Oxford: Oxford University Press, pp. 273–95.
- Sosa, Ernest (2007). *A Virtue Epistemology*. Oxford: Oxford University Press.
- van Inwagen, Peter (1996). "It Is Wrong, Everywhere, Always, and for Anyone, to Believe Anything upon Insufficient Evidence." In Jeff Jordan and Daniel Howard-Snyder (eds.), *Faith, Freedom, and Rationality: Philosophy of Religion Today*. London: Rowman and Littlefield, pp. 137–53.
- Wedgwood, Ralph (2007). *The Nature of Normativity*. Oxford: Oxford University Press.
- White, Roger (2005). "Epistemic Permissiveness." In *Philosophical Perspectives*, vol. 19, *Epistemology*. Oxford: Blackwell, pp. 445–59.
- Williamson, Timothy (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Williamson, Timothy (2007) *The Philosophy of Philosophy*. Oxford: Blackwell.