

Disagreement and Evidential Attenuation

MARIA LASONEN-AARNIO
University of Michigan

Abstract

What sort of doxastic response is rational to learning that one disagrees with an epistemic peer who has evaluated the same evidence? I argue that even weak general recommendations run the risk of being incompatible with a pair of real epistemic phenomena, what I call *evidential attenuation* and *evidential amplification*. I focus on a popular and intuitive view of disagreement, the equal weight view. I take it to state that in cases of peer disagreement, a subject ought to end up equally confident that her own opinion is correct as that the opinion of her peer is. I say why we should regard the equal weight view as a synchronic constraint on (prior) credence functions. I then spell out a trilemma for the view: it violates what are intuitively correct updates (also leading to violations of conditionalisation), it poses implausible restrictions on prior credence functions, or it is non-substantive. The sorts of reasons why the equal weight view fails apply to other views as well: there is no blanket answer to the question of how a subject should adjust her opinions in cases of peer disagreement.

1. Blanket Views of Disagreement

What sort of doxastic response is rational to learning that one disagrees with an *epistemic peer* who has evaluated the same evidence? In particular, how should a subject adjust her opinion on the matter under dispute, and how confident should she be that her own opinion (as opposed to the opinion of her peer) is correct? Almost all views of disagreement put forth in the recent literature offer some general recommendations, recommendations such as “the subject ought to adjust her opinion at least a little bit in the direction of her peer”, and “the subject ought to be equally confident that she made a mistake as that her peer did”. I what follows I will argue that views that make even weak recommendations run the risk of being either non-substantive or false. In particular, blanket recommendations about how subjects should adjust their opinions in cases of disagreement have ignored what I call the phenomena of *evidential attenuation* and *evidential amplification*.

Though I want to draw a general lesson bearing on most—if not all—views put forth in the debate, the discussion will be centered around a popular view that has considerable intuitive pull, the *equal weight view*. I take the view to say that upon learning that she disagrees with an epistemic peer (and learning nothing else), a subject ought to be equally confident that her own opinion is correct as that the

opinion of her peer is. This, I take it, is what it is to attach equal weights to both opinions.¹ My main aim will be to spell out a trilemma for the view: The first horn is that conditionalisation—and what I argue is the correct way of revising one's opinions in certain peer disagreement cases—is violated. The second is that the view ends up posing unmotivated and implausible constraints on prior credence functions, constraints that insulate higher-order probabilities from first-order ones in a highly problematic manner. The third horn is that the view collapses into the recommendation that subjects conditionalise on their evidence. Hence, the equal weight view is either false or non-substantive. And any blanket view of disagreement is in danger of facing the same predicament.

Here is the plan. In §2, I briefly discuss what the equal weight view is. I examine and question the assumption, taken by many for granted, that assigning equal weights to two opinions entails “splitting the difference” between them. Indeed, numerous criticisms of the view, including those pointing to its putative incompatibility with a Bayesian framework, assume this entailment.² I also say why we should think of the equal weight view as a synchronic constraint on prior credence functions. In §3 I spell out the trilemma for the equal weight view. In §4 I reply to some objections, arguing that these merely reinforce the trilemma. In §5 I discuss a line of thought leading to the conclusion that cases of peer disagreement call for an updating procedure that is an alternative to conditionalisation. I say why the alternative updating procedure proposed does not evade my argument. Before concluding, in §6 I say a bit more about the phenomena that spell trouble for blanket views of disagreement, evidential attenuation and evidential amplification.

2. Assigning Equal Weights and Splitting the Difference

Here is the kind of peer disagreement case that I will focus on. Suzy knows that she and her friend, Ned, are about to evaluate a common body of evidence, and to form an opinion concerning the question of whether p based on that evidence.³ Suzy thinks that she is a fairly good judge, and she regards Ned as her epistemic peer: she thinks that Ned is as likely to form a correct opinion based on evaluating the common evidence as she is. Suzy goes on to evaluate the common evidence. In fact, she responds to the evidence in an ideally rational manner. (Below I say more about what such ideal rationality consists in.) She then learns that Ned disagrees with her,⁴ but she acquires no additional evidence about the circumstances of disagreement.

According to the equal weight view, Suzy ought to end up assigning equal weights to both opinions. I take this to mean that she ought to be equally confident that her opinion was correct as that Ned's opinion was correct. Similarly, she ought to be equally confident that she made a mistake as that Ned did.⁵ The view can be further generalised. In disagreement cases a subject's opinions about how likely she is to have gotten it right ought to in some sense be independent of her own evaluation of the common evidence. Instead, she ought to be guided by opinions that she held before acquiring the relevant evidence and learning about the disagreement and its circumstances. If, for instance, Suzy regarded Ned as twice as likely to get it right as herself, then upon learning that they disagree, she ought to regard Ned as twice as

likely to have gotten it right. Or, if she regarded Ned as equally likely to get it right in circumstances in which she has drunk a bottle of wine and Ned is completely sober, then upon learning that they disagree and are in such circumstances, Suzy should regard both subjects as equally likely to have gotten it right. That is the rough idea.

Defenders of equal weight -style views typically also claim, not at all implausibly, that in circumstances involving disagreement with an epistemic peer, a subject ought to adjust her opinion in the direction of that of her peer. In effect, many assume that if a subject assigns a weight of 0.5 to the two opinion, then she ought to at least come close to “splitting the difference” between them, adopting an opinion that is a straightforward average of the two.⁶ So, for instance, if Suzy believed p and Ned believed $\sim p$, then Suzy ought to now suspend judgment in p . If they took finer-grained attitudes, Ned being 0.2 confident in p and Suzy being 0.8 confident in p , then she ought to now be 0.5 confident in p . Indeed, most attacks on the equal weight view have focused on the implausible consequences of splitting the difference.⁷ But why think that assigning equal weights entails splitting the difference in the first place? Before spelling out my trilemma for the equal weight view, let me briefly mention two diagnoses of why the two are so easily conflated.

First, “correct opinion” is ambiguous. On one reading, a correct opinion is just a *true* opinion. Assume that I learn that whereas I believe p , my friend believes $\sim p$. I assign equal weights to both opinions: I think that believing p is equally likely to be correct—i.e., a belief in a truth—as believing $\sim p$ is. Assuming that regarding p and $\sim p$ as equally likely to be true entails suspending judgment in p , it follows that I suspend judgment in p , thereby splitting the difference between the two opinions. Indeed, in the kinds of cases often used by its proponents to motivate the equal weight view, “correct” can be read as “true”, and “incorrect” or “mistaken” as “false”.⁸ However, when proponents of conciliatory views of disagreement such as the equal weight view speak of a given credence or degree of confidence in a proposition being correct, they have in mind another reading of “correct”, which is *being appropriate or reasonable given the evidence*. Opinions that are correct in this sense reflect one’s evidence, not the truth-value of the relevant proposition. But given this reading of “correct”, there is no straightforward entailment between regarding two opinions as equally likely to be correct and splitting the difference between them.

I suspect that thinking otherwise results from implicitly accepting a principle tying together a subject’s credence in a proposition and her credence in its probability on the evidence, a principle that is analogous to Lewis’s Principal Principle, which ties together a subject’s credence in a proposition and her credence in its chance. The *Evidential Expectation* principle says that a subject’s credence in a proposition ought to equal her expectation of its probability on the evidence, or its evidential probability—that is, a subject’s credence ought to equal her expectation of the correct credence (in the second sense of “correct” discussed above).⁹ Let P_S be a subject’s credence function at a time t , and P_E be the evidential probability function for that subject at t :

Evidential Expectation

$$P_S(p) = \sum_{i=1}^n P_S(P_E(p) = r_i) \times r_i$$

Assume that Suzy is equally confident that her credence of 0.8 in p is correct (i.e. equals the evidential probability of p) as that Ned's credence of 0.2 in p is correct, and that she is certain that one of them has gotten it right. Then, *Evidential Expectation* entails that Suzy ought to split the difference between the two opinions, assigning a credence of 0.5 to p .¹⁰

I for one am a fan of the dictum that one's beliefs should be proportioned to the evidence. I take this to mean that a subject ought to assign a credence to a proposition that reflects the degree to which it is supported by the evidence, or its evidential probability. But then, *Evidential Expectation* will only hold if evidential probabilities themselves obey *Expectation*:

Expectation

$$P(p) = \sum_{i=1}^n P(P(p) = r_i) \times r_i$$

To say that *Expectation* holds for any kind of probability is to make substantial assumptions about it. A simple way of guaranteeing its truth is to assume that there is never any uncertainty about higher-order probabilities, i.e., that probabilities are always luminous: if $P(p) = r$, then $P(P(p) = r) = 1$.¹¹ Applied to evidential probabilities, this embeds a very strong assumption about evidence: whenever the evidential probability (for a subject and time) of a proposition p is r , the evidential probability that the probability of p on the subject's evidence is r is 1.¹² Of course, this isn't the only way of guaranteeing *Expectation*: one might, instead, simply assume that $P(p | P(p) = r_i) = r_i$. However, as far as I can see, *Expectation*—and hence, any principle that entails it—fails for the same types of reasons as the assumption that evidential probabilities are luminous.¹³ To say the least, anyone relying on the principle should be prepared to say why such anti-luminosity arguments fail.

I have mentioned two diagnoses of why one might be tempted to take assigning equal weights to two views to entail splitting the difference between them: first, an ambiguity in “correct opinion” and second, an implicit reliance on *Evidential Expectation*. To say the least, the connection between assigning equal weights and splitting the difference is far from clear. In what follows, I will argue that even the claim that one ought to assign equal weights in cases of peer disagreement is highly problematic. Before spelling out the trilemma for the equal weight view, let me say a few words about the kind of framework I will be operating in: first, about my assumptions regarding rational updating and second, about why I will be operating within what could be characterised as an objective Bayesian framework.

As a default starting point, I will assume that when a subject learns some proposition E , and learns nothing else, she should take E into account by conditionalising on it.¹⁴ Someone might worry that conditionalisation is inadequate for taking certain types of higher-order evidence into account, evidence that one disagrees with a peer being a case in point.¹⁵ In §5 I will have more to say about this, and about whether an alternative updating procedure can avoid the trilemma I sketch. But at any rate, showing that the equal weight view leads to violations of conditionalisation would be an interesting result, especially if, as I shall argue, conditionalisation yields the intuitively correct updates in certain problem cases for the view. Indeed, proponents of the equal weight view have not initially put it forth as a view on which cases of peer disagreement call for updating by a procedure that is an alternative to conditionalisation.¹⁶ But neither has it been put forth as a mere recommendation to conditionalise, or as a hypothesis about the kinds of updates that conditionalisation yields in cases of peer disagreement. If the view cannot be thought of as imposing a constraint on how a rational subject's credences evolve, then presumably, it should be thought of as posing a synchronic constraint on what credences it is rational to have at any one time—and ultimately, on the prior credence function. If this is right, then we should think of the equal weight view as analogous to the *Principal Principle* and *Reflection*. It will say something along the following lines: if you regard another subject as a peer, then your credence function ought to satisfy certain further constraints, constraints guaranteeing that if you conditionalise on the information that you disagree (and have no relevant information about the circumstances of disagreement), then you will end up assigning equal weights to the two opinions.

As was remarked above, according to those who hold conciliatory views such as the equal weight view, learning that I disagree with a peer gives me evidence that I have misevaluated my evidence. But it is somewhat difficult to make good on such a thought within a thoroughly subjective Bayesian framework. First, the kind of misevaluation of evidence at issue seems more substantial than merely failing to meet the somewhat mechanical constraints imposed by a subjective framework, such as obeying the probability axioms and having been arrived at by some form of conditionalisation. Rather, it consists in failing to track the degree to which one's evidence objectively supports the relevant proposition. Second, if pretty much any probabilistically coherent prior function will do, it is difficult to see why learning of a disagreement with a peer provides evidence that one has committed some sort of error. After all, without any reason to think that my friend has a prior credence function largely similar to my own, disagreement is precisely what I should expect!¹⁷

Hence, making sense of the kind of misevaluation of evidence under issue seems to require imposing constraints on rational credence functions that go beyond those proposed by subjective Bayesians. However, it doesn't require the assumption that there is always a unique degree to which a body of evidence supports a proposition and hence, that there is no permissiveness as to which credences are rational. Nevertheless, to simplify the discussion below I will speak of "the correct credence" in a proposition, and assume that a subject is certain that, conditional on disagreeing

with her peer, both of their opinions cannot be correct or rational. I don't think my case essentially rests on an assumption of uniqueness. Besides, even those who hold more permissive views can concede that there are possible peer disagreement cases that obey the uniqueness assumptions I make.

I will now spell out the trilemma for the equal weight view: either it (i) violates conditionalisation, and what I take to be intuitively correct updates, (ii) imposes implausible constraints on prior credence functions, or (iii) is nonsubstantial, collapsing into the recommendation that subjects ought to conditionalise on evidence about the disagreement.

3. The Trilemma

What does a subject learn when she learns that she disagrees with an epistemic peer—does she learn just that they disagree, or does she learn a logically stronger proposition stating exactly how they disagree? For now I will assume the latter: a subject A learns of a peer disagreement by learning a proposition specifying exactly how she disagrees with her peer B, a proposition of the form $\lceil P_A(p) = r \text{ and } P_B(p) = r^* \rceil$, where P_A and P_B are the credence functions of A and B, and $r \neq r^*$.¹⁸ Moreover, for now, I won't assume that a subject's credences are luminous to her. Instead, I will assume that a subject simultaneously learns her own credence in the relevant proposition and the credence of her peer (and learns nothing else).¹⁹ In *Appendix II* I ask what kind of constraint is imposed by the equal weight view on prior credence functions within a context that assumes luminosity. But at this point, suffice it to note that my case won't essentially rely on anti-luminosity assumptions.

Let EWV be the following thesis:

EWV

If A regards B as her epistemic peer (regarding whether p), then upon learning only a proposition of the form $\lceil P_A(p) = r \text{ and } P_B(p) = r^* \rceil$, and learning nothing about the circumstances of disagreement, A ought to be equally confident that r is (*or was*) the correct credence in p as that r^* is (*or was*) the correct credence in p .

The thought is that if A regards B as her peer, then for any specific way in which A might learn herself and B to disagree, she ought to regard B's credence as just as likely to have been the correct response to the common evidence as her own. I will proceed by first arguing that EWV is false, and then considering objections to my argument as an argument against the equal weight view, objections that either point to ways in which the equal weight view does not entail EWV, or that question some of the other assumptions I make. I argue that these objections are unsuccessful, and that the kinds of considerations that create trouble for EWV push proponents of the equal weight view into a trilemma.

In so far as A updates by conditionalisation, EWV entails that already prior to learning about her disagreement with B, A ought to think that conditional on any proposition of the form $\lceil P_A(p) = r \text{ and } P_B(p) = r^* \rceil$, r is equally likely to be the correct credence in p as r^* is. Some readers might be suspicious of EWV

at the outset, for isn't regarding a friend as an epistemic peer perfectly compatible with thinking that certain opinions are simply crazy, and cannot be correct, whether those opinions are held by oneself or one's peer?²⁰ But as I will argue, in fact, nothing as dramatic as regarding certain opinions as downright crazy is needed for EWV to fail: given plausible assumptions A might make about the reliability or competence of herself and her peer, and about ways in which their credences are independent, the principle fails whenever A starts out regarding some credences in p as likelier to be correct than others. For instance, if A starts out regarding a credence of 0.2 as likelier to be the correct credence in p than a credence of 0.8 and these further assumptions hold, straightforward conditionalisation on the information that her credence in p is 0.2 and B's credence in p is 0.8 will yield a situation in which A ends up more confident that her credence was correct than that her peer's was. Such cases are counterexamples to EWV, and I will argue that they are also counterexamples to the equal weight view.

In order to spell out the further assumptions needed for my argument, let me describe a toy picture of how subjects form their credences:

God has chosen an ideal, correct credence in p out of n candidates r_1, \dots, r_n . In fact, she chose r_1 . She has painted numbers corresponding to the values r_1, \dots, r_n onto balls, placing them into a bag in such a way as to assure that *most* balls are painted with the value corresponding to the correct credence. Each subject picks out a ball, adopting a credence in p that corresponds to the number written on the ball, before placing the ball back into the bag and passing it to the next subject.

There are two features of this toy picture that I want to focus on. First, each subject is likely, and as likely as other subjects, to form a correct credence in p , no matter what that credence is. I will refer to this assumption as *Global Competence*. Moreover, because the draws are independent, whether one subject X forms the correct credence is independent of whether another subject Y does so. These propositions are also independent conditional on, say, r_1 being the ideal credence. And even more generally: for any credences r_i, r_j , and r_k , conditional on r_i being the correct credence, whether or not Y assigns to p a credence of r_j is independent of whether or not X assigns to p a credence of r_k .²¹ I will call this assumption *Independence*.

Return now to a case of peer disagreement. Assume that at a time t_1 , prior to learning that she disagrees with her peer B, A has a credence distribution over a finite partition of hypotheses about the correct, ideal credence in p (hypotheses such as " r_i is the correct credence in p at t_1 "), and that she is certain that both her own credence and the credence of B are in line with exactly one of these hypotheses. Assume further that A's credence function satisfies *Independence* and *Global Competence*: A regards both herself and B as globally competent in the above sense, and she regards their credences as independent in the above sense. In so far as there are no restrictions on the number of hypotheses about the correct credence in p that A assigns some non-zero credence to, the assumptions made, together with EWV, entail the following constraint on A's credence function²²:

Indifference

$$\forall r \forall r^* P_A(r \text{ is the correct credence in } p) = P_A(r^* \text{ is the correct credence in } p)$$

In other words, prior to learning that she disagrees with B, A must regard all possible credences in p (that is, all those that she thinks might be correct) as equally likely to be correct.

Indifference seems like a wholly unmotivated constraint on A's credences, not in any way justified by treating B as her peer. Even if some form of indifference held with respect to sets of possible outcomes, the principle would be anything but vindicated. Assume, for instance, that I am about to roll a fair die, and an indifference principle tells me to assign a credence of 1/6 to each of the possible outcomes. Let p be the proposition that the outcome will be 1. At least if I am fairly confident that by a principle of indifference I ought to assign equal credence to each possible outcome, I ought to be more confident that 1/6 is the correct credence in p than that 5/6 is. This is perfectly compatible with regarding B as my peer, and thinking that B is equally likely to assign the correct, rational credence to p as I am.

In *Appendix II* I discuss how things look if one assumes that A's own credences are luminous to her. It turns out that even so, an analogue of EWV leads to imposing a strong constraint on how likely A can regard various hypotheses about the ideal credence in p to be. The constraint is not as straightforward as *Indifference*. Rather, it says the following: if A assigns a credence of r_1 to p , then how likely A can regard various hypotheses about the ideal credence in p to be will depend on how likely she thinks B is to assign those credences to p conditional on r_1 being the ideal credence and B failing to assign to p a credence of r_1 . In a simple case in which A thinks that when B goes wrong, he is equally likely to go wrong in any of the possible ways, we get a constraint like *Indifference* applied only to all credences other than the one A herself holds.

The upshot is that whether or not luminosity is assumed, as long as the assumptions made above (in particular, *Global Competence* and *Independence*) hold, the only way to avoid counterexamples to EWV is to impose implausible synchronic constraints on the credence functions of subjects who treat other subjects as epistemic peers. The argument assumed that A updates by conditionalisation, but the lesson I want to draw is not merely that updating on evidence about peer disagreement calls for a procedure other than conditionalisation. Rather, in so far as constraints like *Indifference* are false, EWV threatens what look to be the intuitively correct updates. To see this, consider a variant of the toy picture described above. Assume that there are two candidate correct credences, High and Low. God has chosen the correct credence and made sure that most of the balls in the bag represent the correct credence. Before learning what was painted on the ball picked out by herself or Ned, Suzy is equally confident that she will pick a ball representing the correct credence as that Ned will. If she has no reason to think that God picked High rather than Low, then upon learning that her ball says "High" whereas Ned's ball says "Low", she should be equally confident that her own credence is correct

as that Ned's credence is. But now assume instead that Suzy knew all along that God was likelier to choose High as the correct credence than Low. Then, upon learning that Ned's ball said "Low" and her own ball said "High", Suzy *should* become more confident that her own credence is correct. In fact, this illustrates the phenomenon I refer to below as *evidential attenuation*.

The focus of the discussion above has been on the question of what opinion it is rational for Suzy to adopt regarding whose original credence was correct. But views put forth about peer disagreement also—and even paradigmatically—make claims about what opinion the subject ought to adopt regarding p , the proposition the disagreement is about. Is there anything about the argument above that casts doubt, for instance, on the claim that Suzy ought to adjust her credence in p in the direction of Ned's credence, or even split the difference between the two? Given what was said above, it is a mistake to think that how much Suzy ought to adjust her credence in p is simply a function of how likely she thinks the two competing opinions are (or were) to be correct. But this by no means entails that which credence in p it is rational for Suzy to adopt floats completely free of such matters. A view maintaining that in some cases of peer disagreement I could be almost certain that I am right, and in others almost certain that my peer is right, but that I should (for instance) nevertheless always average out our opinions, is one that I doubt anyone would want to defend.

Before considering objections to my argument given above, it is worth spelling out why the above considerations block an appealing principle that one might think captures the equal weight view in its full generality, a principle I call the *Independence Constraint*.

The Independence Constraint

Recall the thought that in so far as the equal weight view is not to be regarded as putting forth a new principle concerning how a subject ought to update on evidence about disagreement, we should view it as imposing some sort of synchronic constraint on a subject's credence function. The argument above blocks an appealing candidate for the kind of constraint the view might be regarded as imposing. Here is Elga's statement of the equal weight view:

"Upon finding that an advisor disagrees, your probability that you are right should equal your prior conditional probability that you would be right. Prior to what? Prior to your thinking through the disputed issue, and finding out what the advisor thinks of it. Conditional on what? On whatever you have learned about the circumstances of the disagreement."²³

This chimes with remarks made by Christensen about how, upon learning that one disagrees with another subject, the relevant opinions should be *independent of* one's evaluation of the evidence, or about how one should *bracket* the relevant evidence.²⁴

Within a Bayesian context the above remark invites the following interpretation: the relevant credences in a disagreement case should equal the result of updating a prior credence function on—and only on—the information one has about the

disagreement and its circumstances. Let $P_{\text{Suzy}0}$ be Suzy's credence function at a time t_0 before evaluating a body of evidence E and learning that she disagrees with Ned. Assume that Suzy is certain that the two subjects are going to acquire a common body of evidence and form opinions about proposition p based on the evidence at a later time t_1 . Let $P_{\text{Suzy}1}$ be Suzy's credence function at time t_1 , and $P_{\text{Suzy}2}$ be Suzy's credence function at a yet later time t_2 , once she has learnt about the disagreement as well as its circumstances. Similarly for Ned: $P_{\text{Ned}0}$, $P_{\text{Ned}1}$, and $P_{\text{Ned}2}$ are Ned's credence functions at the relevant times. Propositions d and c are as follows:

d : $P_{\text{Suzy}1}(p) = r$ & $P_{\text{Ned}1}(p) = r^*$.

c : The circumstances of disagreement are such-and-such.

In so far as Suzy updates by conditionalisation, $P_{\text{Suzy}2}$ results from conditionalising $P_{\text{Suzy}0}$ on E , d , and c : $P_{\text{Suzy}2}(\cdot) = P_{\text{Suzy}0}(\cdot | E \& d \& c)$. Here, then, is the constraint on priors, inspired by the above remarks by Elga:

The Independence Constraint

$P_{\text{Suzy}0}(x | E \& d \& c) = P_{\text{Suzy}0}(x | d \& c)$, for any relevant proposition x ²⁵

I take the relevant propositions to be (a) propositions about which credence is correct (the proposition that r is the correct credence to assign to p based on the original common evidence, and the proposition that r^* is the correct credence to assign to p based on this evidence), as well as (b) the proposition p itself to be evaluated. The thought is that as far as these propositions go, one's credences in a disagreement case should be what they would have been had one never updated on E in the first place. It's as if one had *only* learnt d and c . What we have here is an independence constraint: the relevant propositions are independent of the evidence E conditional on a certain kind of disagreement situation obtaining. Another way of putting the point would be to say that certain *judgments screen evidence*: as far as the relevant propositions go, the original evidence is screened off by propositions about the disagreement and its circumstances.²⁶

At first sight, this may look like a promising way of constructing a view that is substantive but doesn't violate conditionalisation. But unfortunately, the kinds of points made above create trouble for the *Independence Constraint*. Let's focus on simple situations in which Suzy learns merely that she disagrees with Ned in a particular way, but doesn't learn anything else about the circumstances of disagreement. Now, at t_0 , Suzy doesn't yet know anything very specific about the evidence E that the two subjects will acquire at t_1 . It is plausible that there will be pairs of credences r and r^* such that at t_0 Suzy regards both as equally likely to be the correct credence to assign to p at t_1 , and regards both as equally likely to be correct conditional on Suzy assigning r and Ned assigning r^* :

$P_{\text{Suzy}0}(r \text{ is the correct credence at } t_1 | P_{\text{Suzy}1}(p) = r \& P_{\text{Ned}1}(p) = r^*) =$
 $P_{\text{Suzy}0}(r^* \text{ is the correct credence at } t_1 | P_{\text{Suzy}1}(p) = r \& P_{\text{Ned}1}(p) = r^*).$

Hence, at t_0 , before evaluating the relevant evidence, Suzy regards herself and Ned as equally likely to get things right conditional on disagreeing in a particular way. Assuming that Suzy conditionalises, the *Independence Constraint* entails that if she acquires a body of evidence E and then learns the proposition $P_{\text{Suzy1}}(p) = r$ & $P_{\text{Ned1}}(p) = r^*$, she should *still* regard r and r^* as equally likely to be the correct credence at t_1 . Hence,

$$\begin{aligned} &P_{\text{Suzy1}}(r \text{ is the correct credence at } t_1 \mid P_{\text{Suzy1}}(p) = r \ \& \ P_{\text{Ned1}}(p) = r^*) = \\ &P_{\text{Suzy1}}(r^* \text{ is the correct credence at } t_1 \mid P_{\text{Suzy1}}(p) = r \ \& \ P_{\text{Ned1}}(p) = r^*) \end{aligned}$$

But by the reasoning given in the Appendices, this poses implausible constraints on Suzy's credence function P_{Suzy1} . For instance, assuming that there are no limitations on the number of hypotheses about the correct, ideal credence that Suzy assigns some non-zero credence to, and assuming that her own credence isn't luminous to her, acquiring E cannot make it rational for Suzy to regard a credence of r as more likely to be ideal than a credence of r^* , or vice versa. We get an analogue of *Indifference*: Suzy must still regard r and r^* as equally likely to be ideal. More generally, the *Independence Constraint* poses what look to be implausible constraints on what credences Suzy can assign to higher-order propositions about the correct, ideal credence in p . As such, it insulates first-order probabilities from higher-order ones in a highly problematic manner.

I will now discuss objections to my argument as an argument against the equal weight view, arguing that considering these objections merely reinforces a trilemma for the view: it either lands into one of the predicaments that EWW faces (imposing implausible constraints such as *Indifference* on subjects' credence functions, or being committed to updates that violate conditionalisation), or else the view is non-substantive, boiling down to the recommendation that subjects conditionalise on their evidence.

4. Objections and Replies

I have argued that EWW is in trouble. But what I have said constitutes an argument against the equal weight view only if it entails EWW, and if the assumptions made (in particular, *Global Competence* and *Independence*) are viable. I want to first discuss the objection that I have misconstrued what is involved in "learning that one disagrees with an epistemic peer". The thought is that the equal weight view was never intended to apply when a subject learns something as specific about the differing opinions as I have assumed, as opposed to just learning that she disagrees with her peer—and hence, that the view does not entail EWW. The second objection is that the equal weight view was never intended to apply when subjects learn something relevant about the circumstances of disagreement, and that is exactly what goes on in the kinds of cases I have discussed. That is, sometimes merely learning how one disagrees with an epistemic peer *is* learning about the circumstances of disagreement. The third objection is that the assumption of *Independence* fails in real-world cases. In the next section I discuss the objection that cases of peer disagreement call for an updating procedure that is an alternative to conditionalisation.

i. Learning that one disagrees with a peer

I have assumed that when a subject A learns that she disagrees with another subject B, she learns not only the proposition that they disagree (i.e. that $P_A(p) \neq P_B(p)$), but a logically stronger proposition stating exactly how they disagree (for some r and r^* such that $r \neq r^*$, she learns that $P_A(p) = r$ and $P_B(p) = r^*$). One might take issue with this claim, insisting that EWV should be revised by restricting it to cases in which a subject learns *just* that she disagrees with another subject:

EWV*

If A regards B as her epistemic peer (regarding whether p), then upon learning only that $P_A(p) \neq P_B(p)$, and learning nothing about the circumstances of disagreement, A ought to be equally confident that her credence is (*or was*) correct as that B's credence is (*or was*) correct.

Now, I have assumed that minimally, if A treats B as her epistemic peer with respect to whether p , then she must regard B as equally likely to get things right, or to have the correct, ideal credence in p :

Equal Likelihood of Correctness

$$P_A(P_A(p) \text{ is correct}) = P_A(P_B(p) \text{ is correct})$$

In effect, this is equivalent to regarding A and B as equally likely to get it right conditional on disagreeing.²⁷

Equal Likelihood of Correctness Conditional on Disagreeing

$$P_A(P_A(p) \text{ is correct} \mid P_A(p) \neq P_B(p)) = P_A(P_B(p) \text{ is correct} \mid P_A(p) \neq P_B(p))$$

Elga, for instance, takes this principle to capture what it is to regard another subject as one's epistemic peer.²⁸

But if the above peer principles are satisfied, and A conditionalises on the proposition that $P_A(p) \neq P_B(p)$, it simply follows that she ends up equally confident that her credence is correct that B's credence is correct. In other words, EWV* is *guaranteed* to be satisfied as long as A updates by conditionalisation. At first sight this sounds like good news for the equal weight view. But I don't think it is. For first, it is far from clear whether the basic intuitions and motivations given for the view are thus restricted to cases in which a subject learns merely that she disagrees with her peer. But more importantly, the above move makes no progress towards giving us a new, substantive view: in so far as *Equal likelihood of correctness* is a necessary condition on treating another subject as one's peer, EWV* is equivalent to saying that one ought to conditionalise on the information that $P_A(p) \neq P_B(p)$. As such, we seem to be left with nothing but the recommendation that one conditionalise on one's evidence. This is just one horn of the trilemma I am posing. Remember that the hope was that the equal weight view would pose some plausible, non-trivial synchronic constraint on a subject's probability function.

ii. *Learning about the circumstances of disagreement*

Defenders of the equal weight view are very explicit that the recommendation to assign equal weights in cases of peer disagreement need not apply if a subject learns something further about the circumstances of disagreement. Clearly, it would be a non-starter to claim that Suzy ought to assign equal weights to both opinions even if she learned, for instance, that Ned has been given a drug that seriously impairs his ability to make the sorts of evaluations called for in their present situation. Suzy's confidence in the correctness of the two opinions ought to be guided by her previous assessment of their respective judging abilities conditional on what she subsequently learns about the conditions under which the judgments were made.²⁹ But she never thought, to start out with, that both parties are equally likely to be correct conditional on disagreeing and Ned having been drugged.

Now, perhaps such a constraint could deal with cases in which the opinion of a peer seems absolutely insane—for instance, cases in which he claims that my share of the restaurant bill is \$450, instead of the \$43 that I arrived at.³⁰ The thought is that sometimes learning about another opinion also involves learning something relevant about the circumstances of disagreement. In the case just described, perhaps I learn that I regard the opinion of my friend as absolutely insane. That is why I don't have to assign equal weights: I never thought that, conditional on us disagreeing and me regarding the opinion of my friend as insane, we are equally likely to get things right. Similarly, assume that Suzy regards a credence of r^* as very unlikely to be the correct credence in p , and subsequently learns that Ned assigns to p a credence of r^* . Isn't this like the restaurant case in that Suzy learns that Ned holds an opinion that she regarded, if not insane, then at least highly likely to be incorrect? Doesn't Suzy learn something relevant about the circumstances of disagreement? Perhaps we don't have a counterexample to the equal weight view after all.

Note first that I have argued that we get counterexamples to EWV (and the equal weight view) even if Suzy *merely* learns a proposition about how she disagrees with Ned. She doesn't, in addition, need to learn, for instance, that she regards Ned's opinion as very unlikely to be correct. So if an appeal to learning about circumstances of disagreement is to work, the claim would have to be that sometimes merely learning how one disagrees with an epistemic peer counts as learning something relevant about the circumstances of disagreement. But perhaps this is fine: proponents of the equal weight view need simply to concede that it's harder *not* to learn anything relevant about the circumstances of disagreement than one initially thought.

However, far from being convinced that such a move can avoid the trilemma sketched above, as far as I can see, it merely reinforces it. First, it is not clear whether it is compatible with Elga's statement of the equal weight view: before evaluating the relevant evidence, Suzy may well have thought that conditional on her assigning to p a credence of r and Ned assigning a credence of r^* , both subjects are equally likely to have gotten it right. Second, the threat that we are dealing with a non-substantive view arises again: whenever the view is in danger of violating conditionalisation, the clause about not learning anything relevant about circumstances of disagreement is being appealed to. In cases in which conditionalisation

does not result in assigning equal weights, it is claimed that the relevant subject learned something about the circumstances of disagreement. What seems to be left is a view that does nothing over and above recommending that one conditionalise on evidence about disagreeing with a peer—together, perhaps, with an *ad hoc*-seeming view about what it is to learn something relevant about the circumstances of disagreement. No progress has been made towards providing a new, interesting constraint on priors that would capture something like the idea that judgments screen evidence.

iii. Independence and real-world cases

The argument I gave above relied on an assumption of independence regarding how subjects form their credences (*Independence*). In particular, conditional on r being the correct credence in p , A's assigning r to p is probabilistically independent of B's assigning r (or any other credence) to p . But one might object that such independence doesn't hold in the real world, since subjects are susceptible to the same biases and errors.³¹ Think, for instance, of the Kahneman and Tversky experiments revealing how certain heuristics lead the majority of subjects to violate simple axioms of probability theory.³² In light of such data, shouldn't one expect *Independence* to fail?

Even if, as a general rule, *Independence* didn't hold in the actual world, it is unclear how this could save the equal weight view. *Independence* was an assumption about a given subject's credence function, not about how things stand in the empirical world. Even if such biases exist, the argument given assumed merely that there is some case in which it is rational for a subject to take her credences to be independent of the credences of her peer in the relevant manner. I take EWV (and the equal weight view) to be a claim about how it is, necessarily, rational for subjects to respond in cases of disagreement, whereas the kinds of biases pointed to are merely contingent. Then, capitalising on failures of *Independence* would require imposing its failure as a constraint on rational credences, at least the credences of subjects who regard others as their epistemic peers. And of course, there is absolutely no guarantee that failures of *Independence* would save the equal weight view from imposing implausible restrictions on priors. If, on the other hand, certain dependence constraints could be imposed as a condition on treating another subject as a peer, constraints that would guarantee assigning equal weights in cases of disagreement, this would once again render the view non-substantive.

In the next section I discuss a final objection, the objection that cases of peer disagreement call for updating by a procedure other than conditionalisation.

5. Revising the Prior Function

Proponents of conciliatory views of peer disagreement often express the thought that upon learning that I disagree with a peer about some question, my new opinions about that question, and about how likely our initial opinions are to be correct, ought to be independent of my own evaluation of the evidence. But assuming that

nothing like the *Independence Constraint* discussed above is viable, one could argue that it is impossible to make sense of such independence if all updating happens by conditionalisation.

Consider a situation in which both Suzy and her peer Ned have conditionalised on some total evidence E , and then learn a proposition stating that their credences in p differ in a specific way. Let P_{Suzy_0} be Suzy's prior credence function (her credence function prior to acquiring evidence E). Assuming that Suzy always conditionalises on new evidence, P_{Suzy_0} fixes how her credences change in response to any evidence she might acquire. For instance, as far as her credence in p is concerned, $P_{\text{Suzy}_0}(p \ \& \ E)$ and $P_{\text{Suzy}_0}(E)$ fix how Suzy responds to evidence E . But for Suzy's credence in p —the credence she forms upon learning that she disagrees with Ned—to be fully independent of her evaluation of the evidence, shouldn't it be independent of these prior credences? This, one might think, shows that evidence about disagreement cannot be taken into account by conditionalisation. For if Suzy conditionalises on d , a proposition about how she disagrees with Ned, then her new credence in p will depend on $P_{\text{Suzy}_0}(p \ \& \ E \ \& \ d)$ and $P_{\text{Suzy}_0}(E \ \& \ d)$, and these prior credences don't seem like they are in any intuitive sense independent of $P_{\text{Suzy}_0}(p \ \& \ E)$ and $P_{\text{Suzy}_0}(E)$.

One might be spurred by this observation to argue that information about a disagreement ought not to be taken into account by conditionalisation. Disagreement provides Suzy with evidence of a very special kind, since it provides evidence undermining her way of responding to evidence E , thereby undermining the correctness of $P_{\text{Suzy}_0}(p \ \& \ E)$ and $P_{\text{Suzy}_0}(E)$. But then, the way in which Suzy takes this new evidence about disagreement into account had better not rely on $P_{\text{Suzy}_0}(p \ \& \ E)$ or $P_{\text{Suzy}_0}(E)$. The more general assumption here is the following: when a subject acquires evidence that a rule or policy she is following is mistaken, incorporating that evidence by using the old rule or policy would be to fail to take the defeating evidence seriously.³³ Suzy's prior function P_{Suzy_0} encodes the policy or rule that guides her in taking new evidence into account. Simply conditionalising on information about the disagreement would be to let her new credences be determined by her priors, which is why Suzy ought not to conditionalise. Rather, she should revise her way of responding to evidence as encoded by her prior credence function, thereby revising the prior function itself.

At the same time, in so far as Suzy is 50% confident that her own credence is correct upon learning that she disagrees with Ned, one might think that Suzy's new credences ought to depend in some way on her priors. Perhaps, then, the right way to think about the required independence is that Suzy's new opinions ought to be *equally dependent* on her own original evaluation of the evidence and on Ned's original evaluation. Consider first a simple case involving disagreement about priors: Suzy learns that whereas her prior credence in p is 0.9, Ned's prior credence is 0.1. In so far as Suzy thinks that she is equally likely to be correct as Ned, and that one of them is bound to be correct, a natural thought would be that Suzy ought to revise her priors by averaging out their prior credences, thus ending up assigning to p a prior credence of 0.5.³⁴

When Suzy and Ned disagree after having evaluated a body of evidence E , matters are not as straightforward, since there are two different credence functions that Suzy might go back to revise. Suzy could either revise her priors, or she could revise the function resulting from updating her priors on E . Consider how Suzy should arrive at her new credence in p . On the first view, Suzy should adopt the average of $P_{\text{Suzy0}}(p \ \& \ E)$ and $P_{\text{Ned0}}(p \ \& \ E)$ as her new prior credence in $p \ \& \ E$, and similarly for her new prior credence in E . These new prior credences will reflect her present expectation of the ideal priors. She should then conditionalise her new prior function on evidence E . This results in “splitting the difference” between her and Ned’s *prior* credences in $p \ \& \ E$ and in E , but it need not lead to splitting the difference between the credences in p they arrived at upon evaluating the evidence E .³⁵ One problem with this suggestion is that Suzy might not know the values of $P_{\text{Suzy0}}(p \ \& \ E)$ and $P_{\text{Ned0}}(p \ \& \ E)$, or of $P_{\text{Suzy0}}(E)$ and $P_{\text{Ned0}}(E)$. On the second view, Suzy will simply average out $P_{\text{Suzy0}}(p|E)$ and $P_{\text{Ned0}}(p|E)$. Hence, she will adopt a new prior credence of p conditional on E that equals her present expectation of the ideal conditional credence. This, in effect, just leads to splitting the difference between Suzy’s credence in p and Ned’s credence in p .

The resulting views raise a plethora of technical worries. For instance, the kinds of updates discussed may not leave Suzy with a probabilistically coherent function. In these cases, how should she recalibrate her other credences to arrive at a probabilistically coherent distribution after the new update? Unlike conditionalisation, the new rules don’t say. And needless to say, those persuaded by diachronic Dutch Book arguments won’t be impressed by the new updates. But let me say why, completely independently of such technical worries, I don’t think that resorting to such views is a way of resisting my argument.

First, it is worth noting that views along these lines resurrect some version of the *Evidential Expectation* principle discussed in §2. The idea is that in a case of peer disagreement, a subject ought to go back to revise her prior function in such a way as to end up with new credences or conditional credences that equal her present expectations of what would, or would have, been ideal. But it seems very hard to motivate such an appeal to expectations without accepting that rational *present* credences should match present expectations of ideal credences—and hence, without accepting *Evidential Expectation*.

But even more importantly, the sort of view proposed doesn’t give a recipe for determining just when a subject ought to assign equal weights to two opinions. Rather, it is a suggestion for what confidence a subject ought to assign to a proposition p *once it is already settled* what weights she assigns to her own opinion in p and the opinion of her peer upon learning that they disagree. But I have questioned precisely whether equal weights should be assigned in all cases of peer disagreement. A blanket recommendation to assign equal weights does not leave room for what I think are very real epistemic phenomena, the phenomena of evidential attenuation and amplification. These phenomena arise because a given piece of evidence can have not only first-order import for how likely a proposition p is, but

also higher-order import bearing on how likely various opinions about p are to be correct. Before concluding, I want to discuss these phenomena in a bit more detail.

6. Evidential Attenuation and Amplification

Assume that, having conditionalised on a common body of evidence E , Suzy forms a credence of r_1 in p . Ned's credence is in fact r_2 , but before learning about the disagreement, Suzy already correctly regards r_1 as likelier than r_2 to be the correct credence in p . I have argued that as long as the assumptions made above hold (in particular, *Global Competence* and *Independence*), upon learning that they disagree Suzy should end up more confident that her credence was correct than that Ned's was.

Cases of this sort manifest the phenomenon I call *evidential attenuation*. Informally, the thought is that sometimes a subject's evidential situation can shield her from the defeating force that certain types of evidence would otherwise have. If it is already likelier on Suzy's evidence that r_1 is the correct credence in p than that r_2 is, then her evidence stubs the defeating force (or at least part of the defeating force) that learning that whereas her credence in p is r_1 , Ned's is r_2 , would otherwise have. The thought is that though Suzy still regards Ned as her peer, she now has evidence to think that she is likelier to have gotten it right conditional on them disagreeing in certain specific ways. The mirror phenomenon is *evidential amplification*, which occurs when one's evidential situation amplifies the defeating force of a piece of evidence. So, for instance, if it is likelier on Suzy's evidence that r_2 is correct than that r_1 is, but she learns that her own credence is r_1 and Ned's is r_2 , then her current evidence amplifies the force that learning about the disagreement has on her confidence in the correctness of her own opinion. In neither case should she end up assigning equal weights to the two opinions.

It is not difficult to think about situations in which, despite regarding myself and my friend as epistemic peers, I have evidence that (partially) stubs the defeating force of learning about certain specific opinions. For instance, assume that on the evidence I have, there is an evil demon at work in the neighbourhood who meddles with peoples' credences in a proposition p by making them assign a confidence of 0.9 to it no matter what the relevant evidence points to. If I then learn that my peer is 0.9 confident and I am 0.1 confident in p , it seems reasonable to assign more weight to my own opinion. Or, assume that I have reason to think that both my peer and I are prone to sometimes radically over-estimate the force of the evidence. Then, learning that one of us is very confident in a proposition, whereas the other is less confident, may suffice to make it reasonable to regard the person with a high degree of confidence as likelier to have committed an error. However, it is not clear whether such cases provide any trouble for the equal weight view, for it seems that in such cases one does have relevant evidence about the circumstances of disagreement—for instance, evidence about the workings of an evil demon.

But as I argued above, appealing to the idea that in any case of evidential amplification or attenuation one has relevant evidence about the circumstances of disagreement is in danger of rendering the equal weight view non-substantial, thereby forcing it into the third horn of the trilemma sketched. Further, there is nothing about the structure of the cases that create trouble for the equal weight view that guarantees the presence of evidence that resembles paradigm examples of evidence about the circumstances of disagreement (evidence about someone having drunk wine, being drugged, etc). The trouble cases I have drawn attention to are ones in which, prior to learning that she disagrees with an epistemic peer, a subject's credence distribution over hypotheses about what the correct credence in a proposition p is don't satisfy constraints such as *Indifference*. But having a certain kind of credence distribution over hypotheses about which credences are correct doesn't require anything like evidence about the possibility of being under the influence of drugs, evil demons, etc. The issue isn't with the content of one's evidence, but with its structure.

The kind of evidential structure I pointed to that (given further assumptions) enables evidential attenuation was one where the evidence has a certain degree of awareness about what it supports. For instance, p is likely on the evidence, but it is also likely on the evidence that p is likely on the evidence. In effect, I would conjecture that such situations are not at all atypical. One explanation for such correlations between first- and higher-order probabilities is that the acquisition of first-order evidence bearing on a proposition p is often accompanied by the acquisition of further evidence that has a higher-order nature in being evidence about the first-order evidence for p . For instance, assume that I read in the New York Times that p is the case. Since I possess no further evidence to the contrary, and have no reason to suspect that the newspaper is not a reliable source on the matter of whether p , it is now likely on my evidence that p . But in the course of acquiring the evidence that the New York Times claims that p , I also became aware of the fact that I justifiably believe and know that the New York Times claims that p —and I also know that the fact that the New York Times claims that p is good evidence for p .

However, the kinds of correlations between first- and higher-order probabilities under discussion don't require being able to separate the first- and higher-order contributions of one's evidential situation. For instance, assume that in the absence of defeaters, having a paradigm experience as of rain suffices to make it likely on my evidence that it is raining. But in having a paradigm perceptual experience as of rain, I am aware of having an experience as of rain. Indeed, my awareness of having that experience seems constitutive of its phenomenal character—perhaps I simply could not have the kind of perceptual evidence I have without being aware that I am having an experience as of rain. And by being aware that I am undergoing a paradigm experience as of rain (and further, perhaps, that there are no defeaters), I seem to thereby have evidence that I have evidence of a perceptual sort that it is raining. On the evidence I have, not only is it highly likely to be raining, but it is also highly likely that a high confidence in rain is rational.

More generally, being in a good enough epistemic position with respect to a proposition p may suffice to put one into a good epistemic position with respect to the proposition that one's evidence supports p . Here is a very rough idea. Assume that a subject's evidence consists in all propositions that she bears some evidential relation R to (a relation such as knowing, justifiably believing, etc.). At least sometimes the relation will iterate: a subject will bear R to a proposition p but also bear R to the proposition that she bears R to p . Moreover, as in the perceptual example described above, sometimes the very epistemic circumstances that enable one to bear R to p will put one into a position to bear R to the proposition that one bears R to p . For instance, sometimes the very circumstances that enable one to know p also enable one to know that one knows p . Assume, then, that a subset of Suzy's evidence bears on the question of whether p , and for each proposition in that subset, she knows that she knows it (or more generally, for each of these propositions, she bears R to the proposition that she bears R to it). Let us concede that it is then part of her evidence that these propositions are part of her evidence. Further, assume that Suzy knows that she doesn't have other evidence that bears on p . If Suzy is knowledgeable about her evidence in this way, then her evidence will have the required sort of awareness about itself: not only will the evidence support p , but it will support the claim that it supports p . It is not at all implausible that we often have such access to our evidence.

I argued above that the fact that it is rational to regard certain hypotheses about the correct, ideal credence in a proposition p as likelier than others already prior to learning that one disagrees with an epistemic peer can suffice to break the symmetry in cases of peer disagreement, making it rational to assign more weight to one opinion than another. The phenomena of evidential attenuation and evidential amplification explain how it can be rational to violate the equal weight view.

Conclusions

My intention has not been to merely observe a conflict between the equal weight view and conditionalisation, but to point out that there are cases in which what seems like the correct update leads to counterexamples for the view. The problem cases arise from paying close attention to higher-order probabilities, or to a subject's credences about how likely various credences are to be correct. Attempts to escape the counterexamples lead either to rendering the equal weight view equivalent to saying that subjects ought to conditionalise on their evidence, or to posing implausible restrictions on prior credence functions. I haven't discussed other views of disagreement. However, the above considerations—in particular, the phenomena of evidential attenuation and amplification—cast serious doubt on even fairly mild recommendations about how a subject ought to weight two opinions in cases of peer disagreement, such as the recommendation that a subject should always give at least *some* weight to the opinion of her peer.³⁶ Sweeping generalisations should not be a substitute for investigating how one's initial evidence and evidence about disagreement play together in individual cases.³⁷

Appendix I

Let $\{P_O(p) = r_1, \dots, P_O(p) = r_n\}$ form a partition of hypotheses about the correct, ideal credence in p . One might worry that there are uncountably many such credences, but to avoid such problems, we need not think about each r_i as a point value; I leave open the possibility that these are intervals. What is important is just that r_1, \dots, r_n are disjoint. Hence, the partition might, for instance, consist of three hypotheses: that p is likely, that p is unlikely, and that p is neither likely nor unlikely. Moreover, assume for simplicity that A is certain that both her and B's credence in p is in line with exactly one of these hypotheses. The only other assumptions I will make are the *Independence* and *Global Competence* assumptions discussed above.

Since we will only be considering credences in one proposition p , I will abbreviate ' $P_O(p) = r_i$ ' as ' O_i ', ' $P_A(p) = r_i$ ' as ' A_i ', and ' $P_B(p) = r_i$ ' as ' B_i '. Also, since the question concerns A's credences, I will write ' P ' instead of ' P_A '. P is the function that A has prior to learning that she disagrees with B, that is, learning a proposition of the form $\lceil A_i \& B_j \rceil$. Then, the assumptions made amount to the following, with ' \vee ' for disjunction, for all i, j, k ,

1. $P(\bigvee_{1 \leq i \leq n} (O_i)) = 1$, for all $i \neq j$, $P(O_i \& O_j) = 0$, and for all i , $P(O_i) > 0$
2. $P(\bigvee_{1 \leq i \leq n} (A_i) \& \bigvee_{1 \leq i \leq n} (B_i)) = 1$, and for all $i \neq j$, $P(A_i \& A_j) = P(B_i \& B_j) = 0$
3. $P(\bigvee_{1 \leq i \leq n} (O_i \& A_i)) = P(\bigvee_{1 \leq i \leq n} (O_i \& B_i)) > 0$
4. $P(A_j | O_i) = P(A_j | O_i \& B_k)$
5. $P(A_i | O_i) = P(A_j | O_j)$ and $P(B_i | O_i) = P(B_j | O_j)$

1. is the assumption that the different hypotheses $\{P_O(p) = r_1, \dots, P_O(p) = r_n\}$ about the correct credence in p form a partition, and that A assigns a non-zero credence to each member of the partition. 2. is the assumption that A is certain that both her own and B's credence is in line with exactly one of these hypotheses. Given 1., 3. entails that A regards herself and B as equally likely to assign the correct, ideal credence to p . *Equal likelihood of correctness* (and hence, *Equal likelihood of correctness conditional on disagreeing*) is satisfied. 4. states *Independence*, and 5. *Global Competence*. Now, I take EWV to entail the following:

6. $P(O_i | A_i \& B_j) = P(O_j | A_i \& B_j)$

One can easily show, without making any further assumptions, that for the case where $n = 2$ (i.e. there are only 2 hypotheses about the ideal credence in p) the assumptions made entail

7. $P(O_i) = P(O_j)$.

In other words, A's (prior) credences must satisfy *Indifference*. Here is the proof:

- 8. $P(V_{1 \leq i \leq n}(O_i \& A_i)) = \sum_{1 \leq i \leq n} P(O_i \& A_i) = \sum_{1 \leq i \leq n} P(A_i | O_i)P(O_i) = \sum_{1 \leq i \leq n} P(A_j | O_j)P(O_i)$ (1., 5.)
- 9. $\sum_{1 \leq i \leq n} P(A_j | O_j)P(O_i) = P(A_j | O_j)\sum_{1 \leq i \leq n} P(O_i) = P(A_j | O_j)$ (1.)
- 10. $P(V_{1 \leq i \leq n}(O_i \& A_i)) = P(A_j | O_j)$ (8., 9.)

Similarly,

- 11. $P(V_{1 \leq i \leq n}(O_i \& B_i)) = P(B_j | O_j)$
- 12. $P(A_j | O_j) = P(B_j | O_j) > 0$ (3., 10., 11.)
- 13. $P(A_i | O_i) = P(B_j | O_j) > 0$ (5., 12.)
- 14. $P(O_i \& A_i \& B_j) = P(O_j \& A_i \& B_j)$ (6.)
- 15. $P(O_i \& A_i \& B_j) = P(A_i | O_i)P(O_i \& B_j)$ (4.)
- 16. $P(O_j \& A_i \& B_j) = P(B_j | O_j)P(O_j \& A_i)$ (4.)
- 17. $P(A_i | O_i)P(O_i \& B_j) = P(B_j | O_j)P(O_j \& A_i)$ (14., 15., 16.)
- 18. $P(O_i \& B_j) = P(O_j \& A_i)$ (13., 17.)
- 19. $P(B_j | O_i)P(O_i) = P(A_i | O_j)P(O_j)$ (18.)

Assume that $n = 2$. Then, for $i \neq j$,

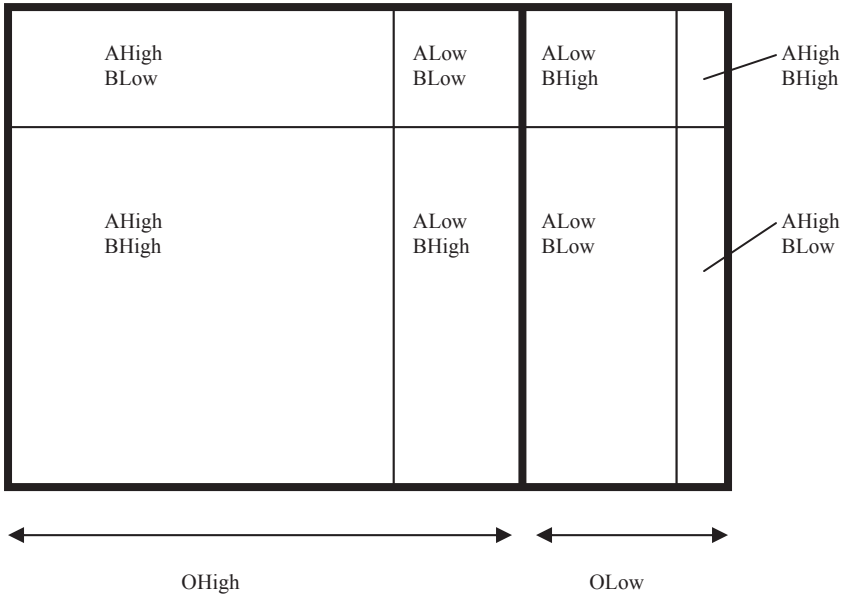
- 20. $P(B_j | O_i) = 1 - P(B_i | O_i) = 1 - P(A_j | O_j) = P(A_i | O_j) > 0$ (2., 5., 12.)

Hence,

- 21. (7.) $P(O_i) = P(O_j)$ (19., 20.)

It follows that when A is certain that the correct credence in p is one of two values then 1.-6. above straightforwardly entail 7.—that is, an instance of *Indifference*—without any further assumptions.

The picture below might help see how failures of *Indifference* and the assumption that subjects update by conditionalisation create what look to be counterexamples to EWV. It represents a case in which the partition of hypotheses about the ideal credence in p only has two members, High and Low. If A conditionalises on the proposition $P_A(p) = \text{High} \& P_B(p) = \text{Low}$ (which I will abbreviate as “AHigh & BLow”), then the only rectangles not ruled out are the two in which this proposition is true. But because the rectangle in which $P_O(p) = \text{High}$ (or OHigh) is true is bigger than that in which OLow is, A will end up more confident of OHigh than OLow and hence, that her own credence is correct. Of course, had she started out more confident of OLow than OHigh, she would have ended up more confident that B’s credence is correct.



Note that focusing on cases in which $n > 2$ doesn't remove the worry that EWV poses implausible constraints on prior credence functions. An example of such a constraint can be seen by considering the fact that 18. entails

$$22. P(O_i) = P(A_i) = P(B_i).^{38}$$

Here is the proof:

$$23. P(A_i \& O_i) = P(B_i \& O_i) \tag{12.}$$

$$24. P(\sim A_i \& O_i) = P(\sim B_i \& O_i) \tag{23.}$$

$$25. P(O_i) = P(O_i \& A_i) + P(O_i \& \sim A_i) = P(O_i \& A_i) + P(O_i \& \sim B_i) \tag{24.}$$

$$26. P(O_i \& \sim B_i) = P(\bigvee_{1 \leq j \leq n, j \neq i} (O_i \& B_j)) = P(\bigvee_{1 \leq j \leq n, j \neq i} (A_i \& O_j)) \tag{2., 18.}$$

$$27. P(O_i) = P(O_i \& A_i) + P(\bigvee_{1 \leq j \leq n, j \neq i} (A_i \& O_j)) = P(\bigvee_{1 \leq j \leq n} (A_i \& O_j)) \tag{25., 26.}$$

$$28. P(A_i) = P(\bigvee_{1 \leq j \leq n} (A_i \& O_j)) \tag{1.}$$

$$29. P(O_i) = P(A_i) \tag{27., 28.}$$

One can similarly show that

$$30. (O_i) = P(B_i),$$

and 22. follows from 29. and 30.

22. entails that A must consider her own and B's credences to track the ideal credences in a very strong way. For instance, if A considers r_1 to be 0.4 likely to be the ideal credence in p , then she must consider both herself and B to be exactly 0.4 likely to assign to p a credence of r_1 .

There are numerous other implausible constraints posed on A's credences even when $n > 2$ that I cannot prove here for reasons of space. For instance, one can show that even if A isn't forced to obey *Indifference*, she cannot regard any of the possible hypotheses about the correct, ideal credence as *much* likelier to be ideal than others. I see no such constraints as rationally imposed by treating another subject as one's epistemic peer.

Appendix II

A complaint one might have about the results in *Appendix I* is that the assumptions made entail that A's own credences are not luminous to her: she learns what her own credence in p is at the same time as learning what B's credence is. I doubt that proponents of the equal weight view would want to rest the viability of their position on an assumption of luminosity. But it is useful to see that a claim along the lines of EWV imposes strong constraints on priors even assuming luminosity. In the context of luminosity, the new version of EWV will be the following:

EWV_L

If A regards B as her epistemic peer (regarding whether p), and A is certain that $P_A(p) = r_j$, then upon learning only a proposition of the form $[P_B(p) = r_i]$, and learning nothing about the circumstances of disagreement, A ought to be equally confident that r_i is (*or was*) the correct credence in p as that r_j is (*or was*) the correct credence in p .

I will make the following assumptions:

1. $P(V_{1 \leq i \leq n}(O_i)) = 1$, for all $i \neq j$, $P(O_i \& O_j) = 0$, and for all i , $P(O_i) > 0$
2. $P(V_{1 \leq i \leq n}(B_i)) = 1$, and for all $i \neq j$, $P(B_i \& B_j) = 0$
3. $P(A1) = 1$
4. $P(V_{1 \leq i \leq n}(O_i \& A_i)) = P(V_{1 \leq i \leq n}(O_i \& B_i)) = P(O1)$
5. $P(B_i | O_i) = P(B_j | O_j)$

1. is as in *Appendix I*: $\{P_O(p) = r_1, \dots, P_O(p) = r_n\}$ form a partition of hypotheses about the correct, ideal credence in p , and A assigns a non-zero credence to each of the hypotheses. 2. states that A is certain that B's credence is in line with exactly one of these hypotheses. 3. captures the luminosity assumption: I am assuming that A is certain that her own credence is in line with the hypothesis $P_O(p) = r_1$, or O1. Note that because of the luminosity assumption, A's credence that her own credence in p is correct or ideal equals her credence in O1. 4. captures *Equal likelihood of correctness*. Note that because A's credences are luminous to her, the global competence assumption cannot be made regarding A: conditional on O1, A is certain to get it right, but conditional on any other hypothesis O_i , A is certain to get it wrong. Nevertheless, I will assume that A regards B as globally competent: conditional on any of the hypotheses about the correct credence in p obtaining, A regards B as equally likely to assign the correct credence to p . This is what 5. says. I take EWV_L to entail the following:

$$6. P(O1 | Bi) = P(Oi | Bi)$$

For each r_i , let Δ_i be how likely B is to assign a credence of r_i to p conditional on r_1 being ideal and B failing to assign r_1 to p :

$$7. P(Bi | O1 \& \sim B1) = \Delta_i$$

Then,

$$8. P(O1 \& Bi) = \Delta_i \times P(O1 \& \sim B1), \quad \text{for any } i \neq 1.^{39} \quad (2., 7.)$$

Moreover, because it is certain that B assigns to p one of $\{r_1, \dots, r_n\}$,

$$9. \Delta_1 + \dots + \Delta_n = 1.^{40}$$

Then,

$$10. P(O1 \& Bi) = P(Oi \& Bi) \quad (6.)$$

$$11. P(O1 \& \sim B1) = P(O1) - P(O1)^2 \quad (1., 4., 5.)^{41}$$

$$12. P(O1 \& Bi) = \Delta_i \times (P(O1) - P(O1)^2), \quad \text{for any } i \neq 1. \quad (8., 11.)$$

$$13. P(Oi \& Bi) = P(Bi | Oi) \times P(Oi) = P(O1) \times P(Oi) \quad (1., 4., 5.)$$

$$14. \Delta_i \times (P(O1) - P(O1)^2) = P(O1) \times P(Oi), \quad \text{for any } i \neq 1. \quad (10., 12., 13.)$$

Beautifying this a bit,

$$15. P(Oi) = \Delta_i \times (1 - P(O1)), \quad \text{for any } i \neq 1.$$

Now, I take it that none of the assumptions made are incompatible with A treating B as her peer. In effect, in so far as *Equal likelihood of correctness* captures what it is to treat another subject as a peer, the assumptions entail that A treats B as her peer.

There is one special case in which 15. is trivially easy to satisfy, namely, a case in which there are only two hypotheses about the ideal credence in p . Call these High and Low, and assume that A's own credence is in line with High. In such a context, in effect the constraint requires merely that A regards the hypotheses that High is the ideal credence in p and that Low is the ideal credence in p as forming a partition—which was one of the assumptions made to start out with. Hence, for cases in which the relevant partition has only two members, 15. ends up not imposing any constraints on A's credence function. EWV_L is satisfied as long as A conditionalises on evidence about B's credence.

However, when there are more than two hypotheses about the correct, ideal credence to which A assigns a non-zero credence, 15. imposes a constraint that is anything but trivial. It restricts A's prior function in the following way: how likely A considers any of the hypotheses $\{P_O(p) = r_2, \dots, P_O(p) = r_n\}$ about the correct credence in p is fixed by how likely she thinks B is to assign various of the

candidate correct credences to p conditional on r_1 being correct but B failing to assign a credence of r_1 . Again, what we have is a constraint on the credences A can assign to various hypotheses about the correct credences in p . For instance, if A thinks that B is equally likely to go wrong in any of the possible ways (conditional on r_1 being ideal but B failing to assign a credence of r_1 to p), then $\Delta_2 = \dots = \Delta_n$. It follows that A must be indifferent among all hypotheses other than $P_O(p) = r_1$, regarding any of the other candidate credences $\{r_2, \dots, r_n\}$ as equally likely to be correct. This yields a constraint very similar to *Indifference*.

Notes

¹ Cf. Elga (2007).

² See Shogenji (2007) and Jehle and Fitelson (2009). White (2009), however, argues against the compatibility of the equal weight view and Bayesianism without assuming any such entailment.

³ It may be objected that no two subjects ever have exactly the same evidence (for instance, it may be that by knowing that I am Maria, I know something that my friend is simply not in a position to know). However, the assumption that two subjects can share the same evidence relevant for some proposition p doesn't seem all that unrealistic. Besides, all that is really needed to get the dialectic going is the assumption that it is possible for two subjects to have, and know that they have, bodies of evidence that are relevantly similar as far as the question of whether p is concerned.

⁴ Exactly what this involves will be one of the main issues raised below: does Suzy learn just the proposition that Ned and her disagree, or does she learn a more specific, logically stronger proposition about how her own opinion and that of Ned differ?

⁵ Elga (2007) gives the clearest statement of the view. See also Christensen (2007: 197) and Feldman (2005, 2006). Elga (2007: 488) writes: "Suppose that before evaluating a claim, you think that you and your friend are equally likely to evaluate it correctly. When you find out that your friend disagrees with your verdict, how likely should you think it that you are correct? The equal weight view says: 50%".

⁶ See, for instance, Elga (2007: 489) and Kelly (2010).

⁷ For instance, Kelly (2010). Several attacks on the equal weight view have focused on a tension between the requirement that a subject ought to split the difference and fundamental Bayesian assumptions. See Shogenji (2007) and Jehle and Fitelson (2009).

⁸ Take, for instance, Elga's (2007: 486) horse race case or Christensen's (2007: 193) restaurant case. In the first case, being correct is judging the winning horse to win. In the second, being correct is coming up with the right sum.

⁹ Christensen (2010b) discusses a principle he calls *Rational Reflection*, which Entails the *Evidential Expectation Principle* assuming that we take the maximally rational credences that Christensen talks about to be probabilities on one's evidence.

¹⁰ It is worth noting, however, that Suzy should *not* end up assigning a 0.5 credence to 0.2 (0.8) being a correct response to *the total evidence she now has*, as opposed to *E*, the original evidence she evaluated. Assume for simplicity that Suzy is convinced that the equal weight view is correct, and that in situations of disagreement, she should move her opinion in the direction of the opinion of her peer. Once she learns that she disagrees with Ned, she should be convinced that neither 0.2 nor 0.8 is *presently* the correct credence to assign to p . Rather, what the equal weight view must be taken to say is that upon learning that she disagrees with Ned, Suzy ought to think that 0.2 and 0.8 are equally likely to have been correct responses to *the original evidence E*. Then, the *Evidential Expectation Principle* can be used to form the subject's "revised" response to *E*, which will be the average of 0.2 and 0.8. I discuss such a view in § 5.

¹¹ Note that I am taking the embedded 'P(p)' as a definite description with narrow scope. As an analogy, take 'P(The person standing in the doorway is male)'. The definite description denotes Bill, but we don't want it to be the case that P(The person standing in the doorway is male) = P(Bill is male). After all, I might be certain that Bill is male, but not that the person standing in the doorway is male. Assuming that we can think of the probability of a proposition p as the measure of accessible worlds in

which p is true, this is to say that when evaluating probabilities of claims involving definite descriptions such as those above, we want to ask the question “What is the measure of accessible worlds w such that the person standing in the doorway at w is male at w ?” and not the question “What is the measure of accessible worlds w such that the person *actually* standing in the doorway is male at w ?”.

¹² Cf. Williamson (2000), pp. 230–237 and 314–315.

¹³ Note that *Expectation* at least requires that certainty is luminous: if $P(p) = 1$, then $P(P(p) = 1) = 1$. Williamson (2008) shows how the basic structure of cases used to construct anti-luminosity arguments—the possibility of constructing sorites series between radically different cases—can be employed to construct arguments against principles like *Expectation*. Assuming that worlds in set W form such a series, the thought is that given any world in W , it is not certain that one is in that world rather than one of its immediate neighbours. Let w be a world in which $P(p)$ attains its maximum value (of all the values it has in worlds in W). If *Expectation* is to hold in w , all of the immediate neighbours of w also have to be ones in which $P(p)$ attains its maximum value. It follows that the value of $P(p)$ must remain constant across *all* the worlds in W . But it is possible to construct sorites series linking two worlds in which the value of $P(p)$ is not the same. Williamson (2008) also generalises the argument to infinite cases.

¹⁴ This is to say that if P_{OLD} is the subject’s old credence function, then upon acquiring evidence E and nothing else (and losing no evidence), her new credence in any proposition p ought to be determined as follows: $P_{\text{NEW}}(p) = P_{\text{OLD}}(p|E)$. However, I won’t need to assume for the present purposes that (strict) conditionalisation is the only way in which the credences of a rational subject can evolve, or even that it is the only way in which a rational subject’s credences can evolve as a response to acquiring new evidence.

¹⁵ Christensen (2010a) is open to the possibility that any undermining evidence of a higher-order nature forces violations of conditionalisation, since taking such evidence into account requires *bracketing* evidence one already has. See also Christensen (2011) and Feldman (2005) for a discussion of higher-order evidence.

¹⁶ For instance, Adam Elga (2007) clearly intends the view to be compatible with conditionalising on information about the disagreement.

¹⁷ This is not to say that evidence about disagreement or agreement can’t have any evidential bearing on whether a subject’s attitude is rational within a subjective framework. For instance, if I am fairly confident that my prior credence function is relevantly similar to that of my friend, learning that we disagree might provide me with evidence that I have failed to conditionalise.

¹⁸ In particular, these are the credence functions that A and B have prior to learning that they disagree. In cases of peer disagreement this is typically not the ultimate prior credence function, since the two subjects have already acquired (at least) a common body of evidence E .

¹⁹ In effect, as I argue elsewhere, I see no in principle difference between how evidence about one’s own credences and evidence about other subjects’ credences ought to be taken into account.

²⁰ Indeed, even proponents of equal weight -like views have wanted to make room for such cases. For instance, Christensen (2007) discusses a variant of the restaurant case in which my friend comes up with an answer that is simply insane.

²¹ This is not to say that whether Y forms a credence of r_1 in p and whether X does so are independent. After all, that Y forms a credence of r_1 makes it likelier that God has chosen r_1 as the ideal credence, thereby making it likelier that X forms a credence of r_1 as well.

²² See Appendix I.

²³ Elga (2007: 490). Elga states that this formulation assumes that the relevant opinions arrived at are all-or-nothing, that is, beliefs in a claim or its negation. He gives an alternative formulation applicable to cases in which this assumption is relaxed. This formulation replaces the first sentence of the above by “Your probability in a given disputed claim should equal your prior conditional probability in that claim”. It is not clear to me why this change is required. For presumably, if anything, how likely you think, after having learnt about the disagreement, that your respective opinions are correct should be guided by your prior assessments of how likely the two of you are to be correct in circumstances of a certain type.

²⁴ See Christensen (2010a, 2011). However, as was remarked above, Christensen is at least sympathetic to a view on which higher-order evidence cannot be taken into account by conditionalisation.

²⁵ Various subtleties must be dealt with: presumably, *E* in itself cannot include information about the circumstances of disagreement, or information that would make it unreasonable for Suzy to continue treating Ned as her peer. Also, sometimes a subject will not learn anything about the circumstances of disagreement. In those cases, we can regard *c* as a necessarily true proposition, giving no new information. Alternatively, we can formulate another constraint stating that the following also holds: $P_{\text{Suzy0}}(x | E \ \& \ d) = P_{\text{Suzy0}}(x | d)$.

²⁶ I heard Brian Weatherson discuss a similar principle in a talk titled “Do judgments screen evidence?”.

²⁷ Because $\{P_A(p) \neq P_B(p), P_A(p) = P_B(p)\}$ form a logical partition, $P_A(P_A(p) \text{ is correct}) = P_A(P_A(p) \text{ is correct} | P_A(p) \neq P_B(p)) + P_A(P_A(p) \text{ is correct} | P_A(p) = P_B(p))$. Similarly for $P_A(P_B(p) \text{ is correct})$. But $P_A(P_A(p) \text{ is correct} | P_A(p) = P_B(p)) = P_A(P_B(p) \text{ is correct} | P_A(p) = P_B(p))$ – both subjects must be equally likely to get it right conditional on agreeing. It follows that $P_A(P_A(p) \text{ is correct}) = P_A(P_B(p) \text{ is correct})$ if and only if $P_A(P_A(p) \text{ is correct} | P_A(p) \neq P_B(p)) = P_A(P_B(p) \text{ is correct} | P_A(p) \neq P_B(p))$.

²⁸ Elga (2007, p. 487, note 21).

²⁹ See, for instance, Elga (2007).

³⁰ See, for instance, Christensen (2007). It may be that the right account of this case is that I was already certain, even before doing the calculation, that my share was within a certain range not including \$450. But we can imagine other cases in which a certain opinion strikes me as insane only once I have evaluated the evidence.

³¹ Thanks to Jim Joyce for pointing my attention to this way of resisting the argument.

³² Kahneman & Tversky (1972).

³³ I argue in Lasonen-Aarnio (*forthcoming*) that this principle is false, though nothing I say below rests on this.

³⁴ Though note also that whether or not Suzy treats Ned as a peer in the first place still depends exclusively on her prior function.

³⁵ The following doesn’t always hold: $\frac{1}{2} \times (P_{\text{Suzy0}}(p|E) + P_{\text{Ned0}}(p|E)) = \frac{1}{2} \times (P_{\text{Suzy0}}(p \ \& \ E) + (P_{\text{Ned0}}(p \ \& \ E)) \div \frac{1}{2} \times (P_{\text{Suzy0}}(E) + P_{\text{Ned0}}(E))$.

³⁶ For instance, though much of what I say is close to the spirit of Tom Kelly’s (2010) “total evidence view”, numerous remarks made by Kelly indicate a view on which all cases of peer disagreement call for at least *some* adjusting of one’s opinion, and *some* decrease in a subject’s confidence that she is correct.

³⁷ I am indebted to Ville Aarnio, Frank Arntzenius, Dave Baker, Gordon Belot, Cian Dorr, David Christensen, Antony Eagle, Adam Elga, Yang-Hui He, Yoaav Isaacs, Jim Joyce, Anna Mah-tani, David Manley, Sarah Moss, Eric Swanson, Teruji Thomas, Rich Thomason, Brian Weatherson, Tim Williamson, and Alastair Wilson.

³⁸ Thanks to Jim Joyce for pointing out this entailment to me.

³⁹ Note that when $i \neq 1$, $P(O1 \ \& \ Bi) = P(O1 \ \& \ Bi \ \& \ \sim B1) = P(O1 \ \& \ \sim B1) \times P(Bi | O1 \ \& \ \sim B1)$.

⁴⁰ Of course, $\Delta_1 = 0$.

⁴¹ i. $P(O1 \ \& \ \sim B1) = P(O1) \times P(\sim B1|O1)$

ii. $P(\sim B1|O1) = 1 - P(B1|O1)$

iii. $P(O1 \ \& \ \sim B1) = P(O1) \times (1 - P(B1|O1))$

iv. $P(B1|O1) = P(O1)$

v. $P(O1 \ \& \ \sim B1) = P(O1) \times (1 - P(O1)) = P(O1) - P(O1)^2$

(i, ii)

(1., 4., 5.)

(iii, iv)

References

- Christensen, D. (2007) “Epistemology of Disagreement: The Good News”, *Philosophical Review* 116: 187–217.
- Christensen, D. (2010a) “Higher-Order Evidence”, *Philosophy and Phenomenological Research* 81(1): 185–215.
- Christensen, D. (2010b) “Rational Reflection”, *Philosophical Perspectives*, 24: 121–140.
- Christensen, D. (2011) “Disagreement, Question-Begging and Epistemic Self-Criticism”, *Philosopher’s Imprint* 11(6): 1–22.
- Elga, A. (2007) “Reflection and Disagreement”, *Noûs* 41: 478–502.

- Elga, A. (2010) "How to Disagree about how to Disagree", in R. Feldman and T. A. Warfield. (eds.), *Disagreement* (Oxford: Oxford University Press), pp. 175–186.
- Feldman, R. (2005) "Respecting the Evidence", *Philosophical Perspectives* 19: 95–119.
- Feldman, R. (2006) "Epistemological Puzzles about Disagreement", in S. Hetherington (ed.) *Epistemology Futures* (New York: Oxford University Press).
- Jehle, D. and B. Fitelson (2009) "What is the 'Equal Weight View'?", *Episteme* 6: 280–293.
- Kahneman, D. & Tversky A. (1972) "Subjective Probability: A Judgment of Representativeness", *Cognitive Psychology* 3(3): 430–454.
- Kelly, T. (2005) "The Epistemic Significance of Disagreement", *Oxford Studies in Epistemology* 1: 167–196.
- Kelly, T. (2010) "Peer Disagreement and Higher-Order Evidence", in R. Feldman and T. A. Warfield. (eds.), *Disagreement* (Oxford: Oxford University Press), pp. 111–174.
- Lasonen-Aarnio, M. (forthcoming) "Higher-Order Evidence and the Limits of Defeat", *Philosophy and Phenomenological Research*
- Shogenji, T. (2007) http://socrates.berkeley.edu/~fitelson/few/few_07/shogenji.pdf
- Weatherson, B. (unpublished) "Do Judgments Screen Evidence?"
- White, R. (2009) "On Treating Oneself and Others as Thermometers", *Episteme* 6(3): 233–250.
- Williamson, T. (2000) *Knowledge and Its Limits*, Oxford: Oxford University Press.
- Williamson, T. (2008) "Why Epistemology Cannot be Operationalised", in Q. Smith (ed.), *Epistemology: New Philosophical Essays* (Oxford: Oxford University Press), pp. 277–300.