

WHAT IS THE “EQUAL WEIGHT VIEW”?

ABSTRACT

In this paper, we investigate various possible (Bayesian) precisifications of the (somewhat vague) statements of “the equal weight view” (EWV) that have appeared in the recent literature on disagreement. We will show that the renditions of (EWV) that immediately suggest themselves are untenable from a Bayesian point of view. In the end, we will propose some tenable (but not necessarily desirable) interpretations of (EWV). Our aim here will not be to defend any particular Bayesian precisification of (EWV), but rather to raise awareness about some of the difficulties inherent in formulating such precisifications.

I. INTRODUCTION

Suppose two agents, S_1 and S_2 , are *epistemic peers* regarding a proposition p : that is, suppose S_1 and S_2 are equally competent, equally impartial, and equally able to evaluate and assess the relevant evidence regarding p (we will call such propositions *p peer-propositions* for S_1 and S_2). After carefully reflecting on the salient evidence for p , suppose S_1 and S_2 discover that they *disagree* about p . For instance, S_1 might believe the defendant is guilty, while S_2 believes the defendant is innocent. Or S_1 might believe that free will and determinism are incompatible, while S_2 believes that the two views are compatible. More generally, S_1 and S_2 might assign *different credences* to p . Examples of peer disagreement (in each of these senses) are common in everyday life, in philosophy, and in many other disciplines.

Question: How should we, if at all, revise our beliefs (regarding p) upon discovering that we disagree with someone we take to be our epistemic peer (regarding p)? Recently several authors have taken this question up and proposed a number of different views. One currently popular and prominent view is the so-called *equal weight view* (EWV) of peer disagreement.¹

In this paper, we will investigate various possible precisifications of the (somewhat vague) notions of “equal weight” that are floating around this literature. We will show that various proposals that immediately suggest themselves are untenable. In the end, we will propose some tenable (but not necessarily desirable) interpretations of “equal weight”. Throughout our discussion, we will assume a (broadly) Bayesian framework. Our aim here is not to defend any particular

Bayesian precisification of EWV, but rather to raise awareness about some of the difficulties inherent in formulating such precisifications.

2. SOME INTUITIONS BEHIND “EQUAL WEIGHT”

Before we get into our investigation of EWV, it will be useful to see what motivates the view in the first place. Consider the following case of peer disagreement from Christensen (2007, 193).

Suppose that five of us go out to dinner. It’s time to pay the check, so the question we’re interested in is how much we each owe. We can all see the bill total clearly, we all agree to give a 20 percent tip, and we further agree to split the whole cost evenly, not worrying over who asked for imported water, or skipped desert, or drank more of the wine. I do the math in my head and become highly confident that our shares are \$43 each. Meanwhile, my friend does the math in her head and becomes highly confident that our shares are \$45 each. How should I react, upon learning of her belief?

According to Christensen, the answer is as follows:

If we set up the case in this way, it seems quite clear that I should lower my confidence that my share is \$43, and raise my confidence that its \$45. In fact, I think (though this is perhaps less obvious) that I should now accord these two hypotheses roughly equal credence.

This passage contains a rather clear statement of the EWV, according to which S_1 and S_2 should assign “roughly equal credence” to p upon learning that they assign different credences to p . We’ll consider various precisifications of this (and other) ideas about EWV below.

Despite the intuitive appeal of the view, proponents of the view have so far failed to give a precise *diachronic rule* for “peer updating”, a rule that would state explicitly what S_1 is to do if she discovers her credal value in p is different from her peer’s credal value in p .

To make things more precise, let $\text{Pr}_i^o(p)$ be the credence S_i assigns to p at t_o , and let $\text{Pr}_i^1(p)$ be the credence S_i assigns to p at t_1 , where $t_1 > t_o$, and between t_o and t_1 , S_1 and S_2 learn that $\text{Pr}_1^o(p) \neq \text{Pr}_2^o(p)$.² What we seek is a rule for (or at least a more precise characterization of) what $\text{Pr}_1^1(p)$ and $\text{Pr}_2^1(p)$ should be in light of S_1 and S_2 learning about their disagreement regarding p at t_o .

Below, we will consider a few different precisifications of an “equal weight rule” EWR for peer updating. Before we discuss the various precisifications, we will lay down some intuitive *constraints* on EWR that have been discussed in the literature on probability aggregation.

Probabilism (P): $\text{Pr}_1^1(\cdot)$ and $\text{Pr}_2^1(\cdot)$ should be probability functions.

Conditionalization (C): $\text{Pr}_1^1(\cdot)$ and $\text{Pr}_2^1(\cdot)$ should respect *conditionalization* as a constraint on the relationship between $\text{Pr}_1^1(\cdot)$ and $\text{Pr}_1^o(\cdot)$, and $\text{Pr}_2^1(\cdot)$ and $\text{Pr}_2^o(\cdot)$. (This will be clarified below.)

Unanimity (U): $\text{Pr}_1^i(\cdot)$ and $\text{Pr}_2^i(\cdot)$ should not force new point-wise disagreements about credence values concerning peer-propositions on which S_1 and S_2 already agree (at t_o).

Agreement (A): $\text{Pr}_1^i(p) = \text{Pr}_2^i(p)$ for all peer propositions p , *i.e.*, S_1 and S_2 should be *in agreement* on all peer propositions p (at t_1).

Irrelevance of Alternatives (IA): $\text{Pr}_1^i(p)$ and $\text{Pr}_2^i(p)$ should each be functions of $\text{Pr}_1^o(p)$ and $\text{Pr}_2^o(p)$ for each peer-proposition p . That is, for each peer-proposition p , $\text{Pr}_1^i(p) = f_1[\text{Pr}_1^o(p), \text{Pr}_2^o(p)]$, and $\text{Pr}_2^i(p) = f_2[\text{Pr}_1^o(p), \text{Pr}_2^o(p)]$ or some functions f_1 and f_2 .

Preservation of Conditional Independencies/Dependencies (PCI): $\text{Pr}_1^i(\cdot)$ and $\text{Pr}_2^i(\cdot)$ should neither reverse initially agreed-upon assessments of conditional independence/dependence (according to $\text{Pr}_1^o(\cdot)$ and $\text{Pr}_2^o(\cdot)$), nor force new disagreements about relations of conditional probabilistic independence/dependence (already agreed upon at t_o) among the set of peer-propositions for S_1 and S_2 .

As mentioned, these conditions aren't new with us. Analogues of these conditions have been discussed extensively in the literature on Bayesian judgement aggregation, and a number of "impossibility results" on various combinations of these conditions have been known since the 60's.³ While the aggregation problem is different than the peer-updating problem, we will see below that they share some common features. First, let's take a closer look at the above conditions.

(P) and (C) are fundamental Bayesian principles. We won't argue for these here, since they are basic theoretical presuppositions of the very framework we are adopting.⁴

We think (U) should be uncontroversial from the point of view of defenders of EWW. The whole idea behind EWW is that we should "minimize" or "reduce" disagreements with peers (on peer-propositions). If we do so by adopting an EWR which is (sometimes) forced to generate new disagreements (on peer-propositions) that weren't there before, then this would undermine the very spirit of EWW.

So, we'll think of these first three conditions ((P), (C), and (U)) as basic *desiderata* for any adequate Bayesian EWW-rule. The next three conditions, on the other hand, will prove to be more controversial (and less sacrosanct) from a Bayesian point of view.

Constraint (A), which is *strictly logically stronger* than constraint (U), also *seems* quite sensible from an EWW point of view. Ideally, an EWW-er wants to *both preserve* existing agreements *and eliminate* existing disagreements on all peer-propositions. However, as we will see toward the end of the paper, it is not clear (on reflection) whether this stronger constraint (A) should be imposed on a *Bayesian* EWW-er.

Constraint (IA) is a standard assumption made in the context of Bayesian strategies for probability *aggregation* (*i.e.*, deciding on a *consensus* probability assignment for a group of Bayesian agents). Whether it should be imposed as a

constraint on EWW-update rules is far less clear. While (IA) may sound plausible, it conflicts with EWW in the presence of another constraint that has been discussed in the literature on probability aggregation, namely, the preservation of conditional independencies/dependencies (PCI) (Wagner 1984).

Constraint (PCI) has been more controversial than the other constraints in the literature on Bayesian aggregation.⁵ Here are some considerations in support of (PCI) from a peer-updating perspective. First, from an *epistemic* point of view, assessments of (in)dependence can reflect *evidential* relationships induced by an agent's credence function (*viz.*, Bayesian confirmation theory; see also Jeffrey (1987)). In such contexts, we think it would be undesirable for EWR to undermine agreed-upon assessments of these important relations. Second, dismissing (PCI) can have undesirable consequences for Bayesian decision theory. Standard Bayesian decision-theoretic resolutions to Newcomb's problem involve some appeal to the fact that, while the presence or absence of the \$1M in the opaque box is *unconditionally* probabilistically dependent on what the agent decides to do, it is probabilistically *independent* of what the agent does, *conditional* on the appropriate causal hypothesis. As a result, in the absence of (PCI), it would be possible for an agent to start out as a two-boxer but end up a one-boxer simply because she disagreed with an epistemic peer on some of the initial probability assignments in a salient representation of Newcomb's problem—even if there was no disagreement about *causal structure* either before or after learning about the credal disagreement. This also strikes us an unacceptable consequence of denying (PCI).

In the aggregation context, (PCI) and (IA) jointly entail that one of the "peers" is actually a *dictator*, in the sense that their credence function is the only acceptable "consensus probability function" (Wagner 1984). An analogous problem will plague EWW in some cases (one of which will be discussed below). Of course, this is clearly in conflict with the spirit of EWW. As a result, an EWW-theorist cannot (in general) accept both (IA) and (PCI).⁶ Ultimately, we will present an EWW-update rule that can always satisfy (PCI), but which does not satisfy (IA). From a probabilistic point of view, we think this makes sense, since (IA) assumes a kind of "locality" that probabilists shouldn't accept. As we'll soon see, probability distributions have a kind of "non-local" or "holistic" character that makes (IA) untenable for a Bayesian EWW-theorist.

We now turn to various Bayesian proposals for precisifying the intuitive characterizations of "equal weight".

3. PRECISIFICATIONS OF "EQUAL WEIGHT"

3.1. Straight Averaging (*a.k.a.* "Splitting the Difference")

One natural way to render equal weight's peer updating rule would be to apply what we call *two-person straight averaging*. On this approach, when S_1 and S_2 discover they

Table 1. A simple two-atomic-proposition (SA)-example.

p	q	$\text{Pr}_1^\circ(\cdot)$	$\text{Pr}_2^\circ(\cdot)$	$\text{Pr}_{\text{SA}}^1(\cdot)$
T	T	0.1	0.55	0.325
T	F	0.2	0.25	0.225
F	T	0.3	0.15	??
F	F	0.4	0.05	??

disagree regarding a peer-proposition p , they should both adopt a new credence for p that is the straight average of their initial credences for p .⁷ More precisely:

Straight Averaging (SA): If S_1 and S_2 find themselves in disagreement regarding a peer-proposition p at t_0 , then:

$$\text{Pr}_1^1(p) = \text{Pr}_2^1(p) = \frac{\text{Pr}_1^\circ(p) + \text{Pr}_2^\circ(p)}{2}$$

From the perspective of equal weight, (SA) has some intuitively desirable properties. Intuitively, (SA) coheres nicely with some informal remarks in recent literature. For instance, Kelly (forthcoming, 12) has us suppose that

at time t_0 , immediately before encountering one another, my credence for H stands at .8 while your credence stands at .2. At time t_1 , you and I meet and compare notes. How, if at all, should we revise our respective opinions? According to The Equal Weight View, you and I should split the difference between our original opinions and each give credence .5 to H .

As stated, however, (SA) is at best *incomplete*; and at worst *synchronically incoherent*. This is because (SA) doesn't say what we should do in cases where changes to *non-peer* propositions are forced (on pain of synchronic incoherence) by averaging the agents' credences on the peer propositions in the space. To see the problem vividly, consider the following simple toy case (Table 1) involving agents S_1 and S_2 who entertain just two "atomic" propositions: p and q .⁸

Let's assume that there are exactly two peer-propositions in this case: $p \& q$ and $\sim p \& \sim q$.⁹ If S_1 and S_2 both follow (SA), then all we know for sure about distributions resulting from (SA) is what we've written under the heading $\text{Pr}_{\text{SA}}^1(\cdot)$ in Table 1. Because neither $\sim p \& q$ nor $\sim p \& \sim q$ are peer-propositions for S_1 and S_2 , (SA) – as stated – implies nothing about what should happen to their credence values at t_1 . Moreover, we cannot just leave the credences of $\sim p \& q$ or $\sim p \& \sim q$ unchanged from t_0 to t_1 . If we were to do that, then both S_1 and S_2 would end up with credence functions that violate (P). This is because a probability assignment must assign probabilities to the four state descriptions in such a way that they sum to exactly 1, and here neither agent's credence function will satisfy this constraint – unless changes are made to the credences of the *non-peer* $\sim p \& q$ and $\sim p \& \sim q$. So, in order to satisfy both (SA) and (P), both S_1 and S_2 must make changes to the credences they assign to *non-peer*-propositions. But *precisely what*

changes should they make? Perhaps this question *need not* be answered by an EWW-rule *per se*. But this example shows that a *conservative* rendition of (SA) – which instructs us to change *only* peer-proposition credences – is *synchronically incoherent*. For this reason, we will include *within* our EWW-rules some (*quasi-conservative*) advice for changing non-peer credences when such changes are mandated on pain of synchronic incoherence. Specifically, we propose adding a "minimal change" clause to (SA) as follows.¹⁰

Straight Averaging + Minimal Change (SAMC): If S_1 and S_2 find themselves disagreeing about a peer-proposition p at t_0 , then

$$\text{Pr}_1^1(p) = \text{Pr}_2^1(p) = \frac{\text{Pr}_1^\circ(p) + \text{Pr}_2^\circ(p)}{2}$$

And if other changes must be made to $\text{Pr}_1^\circ(\cdot)$ and/or $\text{Pr}_2^\circ(\cdot)$ in order to ensure satisfaction of (P), then the other changes should be made so as to minimize the distance¹¹ of $\text{Pr}_1^1(\cdot)$ and $\text{Pr}_2^1(\cdot)$ from the initial distributions $\text{Pr}_1^\circ(\cdot)$ and $\text{Pr}_2^\circ(\cdot)$.

In Table 2 we compute the (SAMC) distributions for our example:

Table 2. A simple two-atomic-proposition (SAMC)-example.

p	q	$\text{Pr}_1^\circ(\cdot)$	$\text{Pr}_2^\circ(\cdot)$	$\text{Pr}_1^1(\cdot)$	$\text{Pr}_2^1(\cdot)$
T	T	0.1	0.55	0.325	0.325
T	F	0.2	0.25	0.225	0.225
F	T	0.3	0.15	0.175	0.275
F	F	0.4	0.05	0.275	0.175

With this additional caveat, of course, (SAMC) is guaranteed to satisfy (P), and in a "quasi-conservative" fashion. However, (SAMC) is not guaranteed to satisfy (C). To see this, we need to clarify the meaning of (C) in the current context. We will use the notation $\text{Pr}_i^{\circ+r}(\cdot)$ to denote the credence function S_i would have were they to learn (exactly) proposition r at (or just after) time t_0 . And we will use the notation $\overline{\text{Pr}}_i^t(p)$ to denote the credence S_i assigns to p as a result of the application of an equal weight updating rule EWR (to their credence in p at time t). For instance, in this context, $\overline{\text{Pr}}_i^t(p)$ will denote the credence S_i assigns to p as a result of the application of the SAMC updating rule (to their credence in p at time t). Now, we're ready to clarify (C):

Conditionalization (C): Suppose p , q , and $p \& q$ are peer-propositions for S_1 and S_2 (at t_0 and t_1), and also that q remains a peer-proposition for S_1 and S_2 (at t_0) on the supposition that p is true. Then, conditionalization imposes the following two constraints:

$$\text{Pr}_i^{\circ+p}(q) = \text{Pr}_i^\circ(q|p) = \frac{\text{Pr}_i^\circ(p \& q)}{\text{Pr}_i^\circ(p)}$$

Table 3. The (SAMC)-example with p and q also peer-propositions.

p	q	$\text{Pr}_1^\circ(\cdot)$	$\text{Pr}_2^\circ(\cdot)$	$\text{Pr}_i^1(\cdot)$
T	T	0.1	0.55	0.325
T	F	0.2	0.25	0.225
F	T	0.3	0.15	0.225
F	F	0.4	0.05	0.225

and

$$\overline{\text{Pr}_i^{\circ+p}}(q) = \overline{\text{Pr}_i^\circ}(q|p) = \frac{\text{Pr}_i^1(q \& p)}{\text{Pr}_i^1(p)}$$

The first constraint in (C) is just the definition of (classical) Bayesian conditionalization itself. The second constraint in (C) is a *commutativity* requirement. What the second constraint says is that it shouldn't matter whether we (a) learn p first, and then do a peer-update or (b) do a peer-update first, and then learn p . That is, the second constraint in (C) requires that the peer-update *commutes with conditionalization*.¹² Given this clarification of (C) in this setting, we can now see that the example depicted in Table 1 will *already* yield a counterexample to (C). We only need to add the assumption that p and q are also peer-propositions for S_1 and S_2 (and that q remains a peer-proposition for S_1 and S_2 at t_o , on the supposition that p is true). Once we add this assumption, (SAMC) forces S_1 and S_2 to share the same distribution $\text{Pr}_i^1(\cdot)$ at t_1 , which is depicted in the final column of Table 3:

As a result, by the first (C)-constraint, we have:

$$\text{Pr}_1^{\circ+p}(q) = \text{Pr}_1^\circ(q|p) = \frac{\text{Pr}_1^\circ(p \& q)}{\text{Pr}_1^\circ(p)} = \frac{0.1}{0.3} = 0.3333$$

$$\text{Pr}_2^{\circ+p}(q) = \text{Pr}_2^\circ(q|p) = \frac{\text{Pr}_2^\circ(p \& q)}{\text{Pr}_2^\circ(p)} = \frac{0.55}{0.8} = 0.6875$$

And applying (SAMC) to these (disagreed-upon, peer) $\text{Pr}_i^{\circ+p}(q)$'s yields:

$$\overline{\text{Pr}_1^{\circ+p}}(q) = \overline{\text{Pr}_2^{\circ+p}}(q) = \overline{\text{Pr}_i^{\circ+p}}(q) = \frac{0.3333 + 0.6875}{2} = 0.5105$$

But this does *not* match what we get when we compute $\overline{\text{Pr}_i^\circ}(q|p)$ directly, by applying the second (C)-constraint, as follows:

$$\overline{\text{Pr}_i^\circ}(q|p) = \frac{\overline{\text{Pr}_i^\circ}(q \& p)}{\overline{\text{Pr}_i^\circ}(p)} = \frac{\text{Pr}_i^1(q \& p)}{\text{Pr}_i^1(p)} = \frac{0.325}{0.55} = 0.5909 \neq 0.5105$$

Therefore, the example depicted in Table 3 is a counterexample to (C) for the (SAMC) updating rule. Moreover, the (PCI) constraint is also violated in this example, since:

$$\text{Pr}_1^\circ(q) = 0.4 > \text{Pr}_1^\circ(q|p) = 0.3333, \text{ and}$$

$$\text{Pr}_2^\circ(q) = 0.7 > \text{Pr}_2^\circ(q|p) = 0.5909, \text{ but}$$

$$\text{Pr}_i^1(q) = 0.55 < \text{Pr}_i^1(q|p) = 0.5909.$$

Thus, (SAMC) also forces a reversal on the initially agreed-upon assessment of S_1 and S_2 that p and q are *negatively* dependent. After the (SAMC) peer-update, they both change their mind about this and come to agree that p and q are *positively* dependent. Since p and q are both peer-propositions in the present example, this is also a counterexample to (PCI) for the (SAMC) updating rule.¹³

Two final notes on (SAMC). First, (SAMC) satisfies (U). Indeed, since S_1 and S_2 will always end up having the same credences on all peer-propositions, (SAMC) satisfies the stronger constraint (A). Second, (SAMC) satisfies (IA), since for each peer-proposition p , the value of the new credence for p is a function (namely, the straight averaging function) of the values of the old credences for p assigned by S_1 and S_2 .

To sum up: because neither of the Straight Averaging rules can always satisfy *both* (P) and (C), neither yields a satisfactory updating rule from a Bayesian perspective. Nonetheless, perhaps there is some way to get "close" to straight averaging, while still respecting these fundamental constraints (and perhaps other constraints as well). We will consider several "approximate" versions of (SA) in the next section.

3.2. "Approximate" Straight Averaging

In the last section, we saw that (SAMC) can lead to unsatisfactory updates. Perhaps straight averaging is not the best way to understand "equal weight" after all. Interestingly, Christensen says that when faced with a disagreeing peer, "I should *come close* to 'splitting the difference' between my friend's initial belief and my own" (2007, 203, emphasis ours). Inspired by this "approximate splitting" intuition, we will now consider three "approximate" renditions of (SAMC), in increasing order of logical strength. Here is the weakest of the three.

Approximate Straight Averaging + Minimal Change₁ (ASAMC₁): If S_1 and S_2 find themselves disagreeing about a peer-proposition p at t_o , then they should each update on p so that:

$$\text{Pr}_i^1(p) \approx \frac{\text{Pr}_1^\circ(p) + \text{Pr}_2^\circ(p)}{2}$$

where $\text{Pr}_i^1(p)$ is *strictly between* $\text{Pr}_1^\circ(p)$ and $\text{Pr}_2^\circ(p)$ ¹⁴ and where the update is done in a way that satisfies (P) and (C). If additional changes must be made (on non-peer

propositions) to $\text{Pr}_1^o(\cdot)$ and/or $\text{Pr}_2^o(\cdot)$ in order to ensure satisfaction of (P) and (C), then the other changes should be made so as to minimize the distance of $\text{Pr}_1^i(\cdot)$ and $\text{Pr}_2^i(\cdot)$ from the initial distribution(s) $\text{Pr}_1^o(\cdot)$ and $\text{Pr}_2^o(\cdot)$, while maintaining satisfaction of (P) and (C).

Rule (ASAMC₁) is the weakest of the three “approximate” (SAMC) rules we will consider, because it only requires that each peer end up “close to the average” on each peer-proposition. This does not require that the peers end up close to *each other*, since approximate equality is not a Euclidean relation (that is, the fact that two numbers a and b are both close to the third number c does not imply that a and b are close to each other, or, more formally, $a \approx c \ \& \ b \approx c \not\Rightarrow a \approx b$). We will consider two strengthenings of (ASAMC₁) below. For now, let’s see how (ASAMC₁) fares on the examples we’ve been discussing.

As it turns out, non-trivial constraints on possible values of ϵ will be *forced* by (ASAMC₁). Consider the example depicted in Table 3. It turns out that the only way to satisfy (ASAMC₁) in this case is if $\epsilon > 1/16$. So, for instance, if we had a threshold of $\epsilon = 0.05$, we would not be able to satisfy (ASAMC₁) in the example depicted in Table 3.¹⁵

In this example, we can also satisfy (PCI), so long as $\epsilon > 1/16$. So, adding (PCI) as an additional constraint to the problem does not make things any worse here. In general (i.e., in all 2-atomic-proposition models¹⁶), this will be the case. That is, we can always add (PCI) as an additional constraint to (ASAMC₁) without imposing additional constraints on possible values of ϵ .¹⁷

Interestingly, (IA) will *not* be satisfied by (ASAMC₁), or any “approximate splitting” rule, for that matter. This is because “approximate splittings” can be achieved in multiple ways for the same pair of initial credence values. As such, there can be no *function*(s) of said credence values that yields the (ASAMC₁)-updated values.¹⁸

Finally, (ASAMC₁) is perhaps *too weak* in any event, since it allows peers to end up with credences that are *not close to each other* on peer-propositions. And the spirit of EWV seems to require that peers end up with credences (on peer-propositions) that are close to *each other*, in addition to being close to the straight average of (i.e., the midpoint between) the initial credence values. That suggests strengthening (ASAMC₁) to require that peers also end up close to *each other*.

This leads to (ASAMC₂), which adds to (ASAMC₁) the requirement that $\text{Pr}_1^i(p) \approx \text{Pr}_2^i(p)$. Because of the nature of “ \approx ”, however, there remains an important *ambiguity* in the statement of (ASAMC₂). Here are two salient ways in which peers might satisfy (ASAMC₂).

1. $\text{Pr}_1^i(p) = \text{Pr}_2^i(p) = \text{Pr}_c^i(p)$. On this reading, which we will label (ASAMC_{2.1}), *agreement* (A) is ensured on each peer-proposition. But because we cannot (always) *exactly* “split the difference” between the two initial credences (on pain of incoherence—as was shown in the sections above), the consensus

value $\text{Pr}_c^i(p)$ will (sometimes) have to be *closer to one of the initial credences than it is to the other*. As a result, *one of the peers will have to make a larger change (or a larger Δ) to their initial credence than the other peer does*. So, while this reading has agents reaching *exact consensus* on all peer-propositions, it does so in a way that may seem untrue to the “equal weight” slogan, since the two peers will have *unequal* “credence- Δ s”. This entails a violation of what we will call “equal credence- Δ s”, or (EC Δ) for short.

2. $\text{Pr}_1^i(p) \approx \text{Pr}_2^i(p)$, but $\text{Pr}_1^i(p)$ and $\text{Pr}_2^i(p)$ *may remain unequal*. On this reading, which we will label (ASAMC_{2.2}), exact consensus *need not* be reached on all peer-propositions. That is, (A) is *not* ensured. But we will further precisify (ASAMC_{2.2}) so as to ensure that each updated credence $\text{Pr}_i^i(p)$ is *equally far from the midpoint* between the initial credences $\text{Pr}_i^o(p)$. In this way, (ASAMC_{2.2}) will always satisfy “equal credence- Δ s” (EC Δ).

We will not take a stand here on which precisification of (ASAMC₂) is a “better” EWV-update rule. We think this will depend on the relative importance of (A) vs. (EC Δ). If one insists on (A) being enforced, then one must give up (EC Δ). On the other hand, if one is willing to live without (A), then one can enforce (EC Δ). The important point for our purposes is that an EWV-er *cannot have both* (A) and (EC Δ). So, defenders of EWV must choose which of these two constraints is more important, from an EWV point of view.

Be that as it may, (ASAMC₁) and both precisifications of (ASAMC₂) have *formal* properties that are very similar. The same constraints on ϵ are forced on all three (ASAMC)’s by (P) and (C). That is, (in the 2-atomic- p case) we don’t get stronger constraints on ϵ imposed by the (ASAMC₂)’s, even though they are (logically) stronger than (ASAMC₁). Also, (PCI) can always be satisfied by any of the three (ASAMC) rules, and its satisfaction won’t (generally) require a larger ϵ than that already required by the satisfaction of the synchronic and diachronic Bayesian coherence constraints (P) and (C).

The bottom line here is that—so long as ϵ is sufficiently large—*all three* (ASAMC)’s can *always* be successfully applied (and with very similar formal constraints on ϵ). The only question will be whether (ASAMC) solutions can be found that are *within some ϵ -tolerance*. As we saw above, even the fundamental Bayesian coherence requirements (P) and (C) will sometimes force ϵ to be non-trivially large in (ASAMC) updates. And by adding additional constraints (above and beyond (PCI)) to an (ASAMC) update, one can force ϵ to be even larger (see fn. 17). We leave it to the defenders of EWV to decide which additional constraints might make sense, and how large ϵ should be allowed to get, in various contexts. The purpose of this paper is merely to raise awareness about some of the difficulties in formulating a precise EWV-update rule that is compatible with basic Bayesian tenets. We conclude with a table summarizing some of the results we have discussed.

Table 4. Summary of properties of our EWV-update rules.

Rule	Can Rule (Always) Satisfy Condition?						
	(P)	(C)	(U)	(A)	(ECΔ)	(IA)	(PCI)
(SA)	NO ¹⁹	NO	YES	YES	YES	YES	NO
(SAMC)	YES	NO	YES	YES	YES	YES	NO
(ASAMC _{2,1})	YES	YES	YES	YES	NO	NO	YES
(ASAMC _{2,2})	YES	YES	YES	NO	YES	NO	YES

REFERENCES

- Christensen, D.** 2007. "Epistemology of Disagreement: The Good News." *Philosophical Review* 119: 187–217.
- Dalkey, N.** 1972. *An Impossibility Theorem for Group Probability Functions*. P-4862. Santa Monica, CA: The Rand Corporation.
- Diaconis, P. and Zabell, S.** 1982. "Updating Subjective Probability." *Journal of the American Statistical Association* 77: 822–30.
- Elga, A.** 2007. "Reflection and Disagreement." *Nous* 41: 478–502.
- Feldman, R.** 2006. "Epistemological Puzzles About Disagreement." In S. Hetherington (ed.), *Epistemology Futures*, pp. 216–36. Oxford: Oxford University Press.
- Feldman, R.** 2007. "Reasonable Religious Disagreement." In L. Antony (ed.), *Philosophers Without God: Meditations on Atheism and the Secular Life*. Oxford: Oxford University Press.
- Fitelson, B.** 2008. "A Decision Procedure for Probability Calculus with Applications." *Review of Symbolic Logic* 1: 111–25.
- Genest, C. and Zidek, J.** 1986. "Combining Probability Distributions: A Critique and an Annotated Bibliography." *Statistical Science* 1.1: 114–35.
- Greaves, H. and Wallace, D.** 2006. "Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility." *Mind* 115: 607–32.
- Jeffrey, R.** 1987. "Indefinite Probability Judgment." *Philosophy of Science* 54: 586–91.
- Jeffrey, R.** 2004. *Subjective Probability: The Real Thing*. Cambridge: Cambridge University Press.
- Joyce, J.** 1998. "A Nonpragmatic Vindication of Probabilism." *Philosophy of Science* 65: 575–603.
- Kelly, T.** Forthcoming. "Peer Disagreement and Higher Order Evidence." In R. Feldman and T. Warfield (eds.), *Disagreement*. Oxford: Oxford University Press.
- Lehrer, K. and Wagner, C.** 1981. *Rational Consensus in Science and Society: A Philosophical and Mathematical Study*. Dordrecht-Boston: Reidel.
- Lehrer, K. and Wagner, C.** 1983. "Probability Amalgamation and the Independence Issue: A Reply to Laddaga." *Synthese* 55: 339–46.
- Loewer, B. and Ladagga, R.** 1985. "Destroying the Consensus." *Synthese* 62: 79–95.
- Shogenji, T.** 2007. "A Conundrum in Bayesian Epistemology of Disagreement." Unpublished manuscript.

- Wagner, C.** 1984. "Aggregating Subjective Probabilities: Some Limitative Theorems." *Notre Dame Journal of Formal Logic* 25: 233–40.
- Wagner, C.** 1985. "On the Formal Properties of Weighted Averaging as a Method of Aggregation." *Synthese* 62: 97–108.
- Wagner, C.** 2002. "Probability Kinematics and Commutativity." *Philosophy of Science* 69: 266–78.

NOTES

- * We thank the participants of the *Sixth Annual Episteme Conference* (especially Tomas Bogardus, Fabrizio Cariani, David Christensen, Stew Cohen, and Rich Feldman) for useful feedback on an earlier version of this paper.
- Proponents of EWV (in various informal flavors) include Feldman (2006, 2007), Elga (2007), and Christensen (2007).
 - In general, peers will learn *more* about "the circumstances of their disagreement" (Elga 2007) than *merely* $\text{Pr}_1^o(p) \neq \text{Pr}_2^o(p)$. We will assume that they also learn the *numerical values* of $\text{Pr}_1^o(p)$ and $\text{Pr}_2^o(p)$. That information will also be required for the sorts of update-rules we'll be considering. We'll remain neutral on what *else* they might learn about their disagreement. But we do think that the EWV idea makes the most sense when the information they learn is restricted to the nature of their credal disagreement *qua credal disagreement*. For instance, they may also learn things about *intrinsic properties* of the credences they assign to θ at t_o (e.g., that they are both "high"), which should (intuitively) *not* be taken into account by an update rule for responding to *disagreement per se*. This is a subtle philosophical issue, which we won't be able to delve into further here.
 - For an excellent survey of these results, see Genest and Zidek (1986).
 - See Greaves and Wallace (2006), Joyce (1998), and Jeffrey (2004). Note: we will only need to assume here that the agents are synchronically and diachronically coherent over very simple languages containing just *two* atomic sentences. As such, the variety of "ideal Bayesian rationality" we will need here is quite minimal.
 - See Loewer and Laddaga (1985) and Wagner (1985) for the debate about (PCI).
 - A closely related "dictatorship" impossibility result follows from (IA) *alone* if it is required to hold not only for the unconditional probabilities $\text{Pr}_1^i(p)$ and $\text{Pr}_2^i(p)$, but also for (all) conditional probabilities $\text{Pr}_1^i(p|\cdot)$ and $\text{Pr}_2^i(p|\cdot)$ (Dalkey 1972).
 - Lehrer & Wagner (1981, 1983) and Shogenji (2007) discuss averaging proposals for aggregation. Kelly (forthcoming) calls this proposal "splitting the difference".
 - For the purposes of this paper, we will only discuss very simple toy models in which there are just two "atomic" (logically independent) propositions in the agents' doxastic spaces. Some of our results can be lifted to larger spaces, but the technical details (constraint satisfaction, etc.) are exponentially complex. Our purpose here is just to give some sense of the difficulties inherent in clarifying EWV. For this purpose, it is best to give the *simplest possible* problematic examples. We leave it up to (Bayesian) EWV-ers to think about more complex/realistic models.
 - This immediately raises questions about "the logic of peer-proposition-hood". For instance, does it follow from the fact that $\lceil p \ \& \ q \rceil$ is a peer-proposition that p and q are also peer-propositions? For the present example to make sense, the answer to this question must be "no". We think this is the right answer. Here's an intuitive

- counter-example to “conjunction-elimination for peer-proposition-hood”, which we owe to David Christensen. You and I could be peers with respect to identifying flying mammals, and also with respect to identifying flying animals in general. But the only mammals I’m really interested in are bats. I don’t really know if people or whales or platypuses are mammals, while you really know your mammals. So if A is “that’s a flying animal” and B is “that’s a mammal” we could be peers on the conjunction $A \& B$, but not on B . Indeed, this looks like a case in which we’re peers on the conjunction $A \& B$ and on A , but not on B (which, plausibly, is also what the structure of our first example here is like). As far as we know, none of the defenders of EWV have discussed this “logic of peer-proposition-hood” issue. Perhaps when they do, some interesting “logic” will be discovered (and this may render some of our present examples otiose). Again, we leave this for the (Bayesian) defenders of EWV to work out. We will assume no general “logical laws” for peer-proposition-hood. We suspect that there are very few general “logical laws” of this kind. Perhaps a Bayesian should say that the “peer-proposition-hood” of one proposition (q) is determined by the “peer-proposition-hood” of another (p) if *the probability of q is a function of the probability of p* . That “law” (which sounds good to us, but which we won’t defend here) is consistent with all of our examples.
- 10 None of the impossibility (or possibility) theorems in this paper depend essentially on what we say about such (P)-forced changes to *non-peer* credences. That is, nothing we say here trades essentially on the specific choice of a “minimal change” forced-non-peer-changes *addendum* to (SA). We choose this merely for simplicity and concreteness (i.e., so that we have an SA-type rule that gives precise numerical advice in all examples, etc.). Other rules/heuristics could be adopted for the forced non-peer changes, and similar results would obtain. As with the “logic of peer-proposition-hood” (see fn. 9), we leave these “forced-non-peer-change heuristics/rules” (which have also not been discussed in the literature, as far as we know) to be worked-out with more generality (and care) by the defenders of EWV-rules.
- 11 We will assume a Euclidean distance metric, i.e., $\sqrt{(p_1 - q_1)^2 + (p_2 - q_2)^2}$. Other metrics could be used (and similar results would obtain). But since we’re adding this “minimal change” addendum to (SA) merely for simplicity and concreteness in our presentation (see fn. 10 above), we won’t fuss too much over this choice. See Diaconis & Zabell (1982) for a fascinating discussion of the connection between “minimal change” (in the present sense) and Bayesian updating.
- 12 Some Bayesian defenders of EWV require that (ideally) the result of an EWV-update should be *equivalent* to a (classical) conditionalization, which conditionalizes “on whatever you [i.e., both of the agents in a symmetric peer case] have learned about the circumstances of the disagreement” (Elga 2007, 490). If that’s right, then both constraints of (C) will follow from the definition of (classical) Bayesian conditionalization, since pairs of (classical) conditionalizations *must* commute. But even if we don’t think of EWV-rules as *equivalent* to some conditionalization, we think (C) should remain a *desideratum* for EWV-updates. We don’t have the space to defend this claim here. But in general, we are sympathetic to commutativity as a requirement for Bayesian updating. See Wagner (2002) for discussion.
- 13 Examples like this have also been discussed in the literature on Bayesian aggregation. See Shogenji (2007) for an in-depth discussion of (C) – and its interactions with conditions (P), (IA), and (PCI) – in the context of Bayesian aggregation.

- 14 We impose this *strict* between-ness requirement so as to rule out *dictatorial* updates, which revert to one of the two initial assignments. We will assume that $a \approx b$ iff $|a - b| < \epsilon$, for some “small” $\epsilon > 0$. For simplicity, we’ll assume that *the same ϵ is adopted for each peer-proposition*, and we won’t take a stand on what “small” means (or whether any of these things are context-sensitive, etc.). As with our “minimal change” caveat (fn. 10), these assumptions about “ \approx ” and “ ϵ ” could be relaxed/changed. Again, we leave such generalizations to the defenders of EWV.
- 15 Moreover, there exist similar examples in which ϵ is forced to be even greater. We have been able to find examples like these in which ϵ is forced to be larger than 0.1. We omit all technical details here, but a companion *Mathematica* notebook for this paper is available for download at <http://fitelson.org/ew.nb> (a PDF version of the notebook is at <http://fitelson.org/ew.nb.pdf>), which verifies all the mathematical claims made in this paper. There, we present a decision procedure for the class of 2-atomic-proposition models discussed here. That decision procedure is derived from a general decision procedure for the probability calculus (called PrSAT), which is described in (Fitelson 2008).
- 16 These sorts of claims become very difficult to verify when more complex spaces are involved (especially the constraints imposed by (PCI)). Again, we leave such generalizations of the present models and results to the defenders of EWV.
- 17 We could further generalize our (ASAMC)-rules, by allowing additional constraints \mathcal{C} to be added into the updating and minimal-change steps. Our *Mathematica* code (see fn. 15) could easily be changed to allow for arbitrary sets of constraints \mathcal{C} (as long as the \mathcal{C} ’s are jointly consistent with (P), (C), and (PCI), of course).
- 18 This is significant, because (IA) is implicated in most (if not all) of the “impossibility theorems” in the aggregation literature (see Genest and Zidek (1986)). By relaxing (IA), “approximate-splitting” EWV-approaches can avoid these impossibility results. As we explain below, they yield some interesting *possibility* results.

David Jehle completed his Ph.D. at Cornell University in January 2009. In his dissertation, he proved various results in Bayesian confirmation theory and applied those results to issues in contemporary epistemology, political science, and legal theory. He is currently seeking employment in law enforcement.

Branden Fitelson is an Associate Professor in the philosophy department at UC-Berkeley. He also is a member of the Group in Logic and the Methodology of Science and the Cognitive Science Core Faculty at Berkeley. Branden has held non-teaching positions at NASA’s Goddard Space Flight Center and (more recently) at Argonne National Laboratory. Before joining the philosophy department at Berkeley, Branden held teaching positions (in Philosophy) at the University of Wisconsin at Madison, the University of Illinois at Urbana-Champaign, Stanford University, and San Jose State University. His current research interests include the philosophical foundations of induction and probability, automated reasoning (what used to be called automated theorem proving), and the psychology of inference.