A Conundrum in Bayesian Epistemology of Disagreement[1]

Tomoji Shogenji

tshogenji@ric.edu
www.ric.edu/tshogenji/

Abstract

The proportional weight view in epistemology of disagreement generalizes the equal weight view and proposes that we assign to the judgments of different people weights that are proportional to their epistemic qualifications. It is known that (under the plausible Context-Free Assumption) if the resulting aggregate degrees of confidence are to constitute a probability function, they must be the weighted arithmetic means of individual degrees of confidence, but aggregation by the weighted arithmetic means violates the Bayesian rule of conditionalization. The double bind entails that the proportional weight view is inconsistent with Bayesianism. The paper explores various ways to respond to this challenge to the proportional weight view.

[1] A precursor of this paper was presented at the Formal Epistemology Workshop held at Carnegie Mellon University in 2007. I would like to thank my commentator David Jehle and other participants of the workshop, especially Branden Fitelson, Isaac Levi, Teddy Seidenfeld, and Carl Wagner for valuable comments and stimulating discussions.

1. Bayesian Epistemology of Disagreement

Epistemology of disagreement addresses the problem of how best to revise (or not to revise) our beliefs in light of disagreement with others, especially when we disagree with our epistemic peers. One's epistemic peer is someone who possesses the same empirical evidence on the issue as one does, and who possesses the same level of relevant epistemic competence as one does. Many discussants of the subject appear sympathetic to the equal weight view—the view that we should assign as much weight to our epistemic peer's judgment as to our own. The main reason is symmetry.[2] There seems to be no epistemically good reason to assign different weights to the judgments of equally qualified people even if one of the people is me—holding onto one's own judgment just because it is one's own seems unreasonable.

One response to the argument from symmetry is that disagreement itself is good reason not to think of the disagreeing person an epistemic peer any more.[3] That would be an appropriate reaction in some cases, for example, when someone abruptly expresses a bizarre opinion. We naturally suspect the person is in confusion and downgrade her epistemic qualifications at least temporarily. However, when the proffered view is clearly the person's considered opinion and we can see where she is coming from, it seems unreasonable to downgrade her otherwise well-established epistemic qualifications just

---

[2] See for example Christensen (2007).

[3] See for example Kelly (2005).

because she doesn't agree with us.[4] The disagreement would be an equally good reason to question *our* epistemic qualifications.

Despite the strong appeal of the argument from symmetry, many discussants also appear uneasy about the apparent consequence of the equal weight view that we need to be agnostic about every controversial issue on which our epistemic peers disagree with us. It may be fine to suspend our beliefs about minutiae of history, but we find disagreement with our epistemic peers on issues close to our heart. Does epistemic rationality demand that we suspend all our deeply held beliefs, ranging from politics to religions, whenever our epistemic peers disagree with us? Some consider it spineless.[5] Can we perhaps agree to disagree while still remaining rational?[6]

This paper examines the epistemological problem of disagreement in probabilistic terms. The probabilistic approach is needed since many disagreements are disagreements in degrees of confidence. For example, some people might think it very likely that there was life on Mars, while others might think it only somewhat likely. In such cases, the disagreement concerns the appropriate degrees of confidence in the proposition, instead of some people believing it while others disbelieving it. One natural way of formulating the equal weight view in probabilistic terms would be "splitting the difference"

---

[4] As Adam Elga (forthcoming) points out, that would make it too easy to demonstrate one's epistemic superiority.

[5] See for example Kelly (2005). Similar points are made by van Inwagen (1996). For some attempts to deflect the charge of spinelessness, see Christensen (2007) and Elga (forthcoming).

[6] Aumann (1976) has shown that we cannot agree to disagree in the following sense: If two people assign the same prior probabilities, and their posterior probabilities for an event are common knowledge (i.e. each knows them, each knows that the other knows them, each knows that the other knows that the other knows them, and so on), then these posterior probabilities are equal. But, of course, many disagreements arise because of different prior probabilities.

(Christensen 2007, p. 203) in the sense of taking the mean of the two degrees of confidence. Here's one implementation of this idea.

We start with the usual Bayesian assumptions that the epistemic subject S_i's degrees of confidence are probabilistically coherent (her credence function C_i is a probability function) and that S_i's degrees of confidence are updated in accordance with the Bayesian rule of conditionalization—i.e. $C_i^{+r}(p) = C_i(p|r) =_{\text{def.}} C_i(p \,\&\, r)/C_i(r)$, where $C_i^{+r}(p)$ is S_i's confidence in *p* after the truth of *r* is known to S_i. Two people S_i and S_j are in disagreement on *p* iff $C_i(p) \neq C_j(p)$. If we let $Q = <q_1, q_2>$ represent S_1's and S_2's epistemic qualifications, where $q_1 + q_2 = 1$, then S_1 and S_2 are epistemic peers iff $q_1 = q_2$. If we split the difference, then the revised degree of confidence $C_1{*}(p)$ should be $½C_1(p) + ½C_2(p)$. The opponents of the equal weight view, on the other hand, would give a bigger weight to her own degree of confidence, $C_1{*}(p) = w_1C_1(p) + w_2C_2(p)$, where $w_1 > w_2$.

There is however another way of understanding the equal weight view, which is to regard both $C_1(p)$ and $C_2(p)$ as permissible degrees of confidence without seeking a unique degree of confidence.[7] This seems sensible in certain cases, e.g. when the disagreeing people are independently engaged in an on-going research. From the perspective of the research community, it is better for each person to try out their own ideas, while keeping their mind open, instead of everybody pursuing the same line of inquiry based on the same degrees of confidence on the competing hypotheses.[8] But the situation is different when we are consumers of information and need to make an

---

[7] See for example Levi (1980).

[8] See for example Kitcher (1990). This point may help alleviate to some extent the concern mentioned earlier that the equal weight view is spineless.

imminent decision. For example, when we decide on a medical treatment, it is better to form a unique degree of confidence to maximize the expected utility, than having two permissible degrees of confidence. The split-the-difference approach, which suggests we take the mean of the two degrees of confidence, seems sensible in such a situation even if the decision-maker is also a researcher and qua researcher she would pursue her own ideas. The slit-the-difference approach also seems sensible when we give advice to consumers of information faced with decisions. It is unreasonable to ignore our epistemic peer's view in advising a layperson even though qua researcher we may well pursue our own ideas. In what follows I will examine the idea of splitting the difference, assuming that the situation is appropriate for this approach. Hereafter I will use the term "the equal weight view" to refer to the split-the-difference version of the equal weight view.

As mentioned already, the opponents of the equal weight view would give more weight to their own view than to their epistemic peer's. This must be distinguished from cases where the disagreeing people are not (and are not considered to be) epistemic peers to begin with, e.g. they do not have the same empirical evidence on the issue, or they do not possess the same level of relevant epistemic competence. The supporters of the equal weight view would not insist on giving the same weight to everyone's judgment in such cases. They would give a bigger weight to an epistemic superior's judgment and a smaller weight to an epistemic inferior's judgment. In fact the natural extension of the equal weight view is to give everyone's judgment a weight that is proportional to their epistemic qualifications. To put this formally, $C_1{*}(p) = q_1C_1(p) + q_2C_2(p)$. The equal weight view is a special case where $q_1 = q_2$. We can extend it to cases involving three or more persons,[9] namely: $C_1^*(p) = q_1C_1(p) + \ldots + q_nC_n(p)$, where $q_1 + \ldots + q_n = 1$.[10] The opponents of the equal weight view would object to the proportional weight view as well and give a bigger weight to one's own judgment than the epistemic qualification warrants.[11]

2. Bayesian Double Bind

The proportional weight view seems reasonable at first glance, but it soon becomes clear that the view, as it stands, is unacceptable. I have already stated the Bayesian assumptions that C_i is a probability function, and that S_i updates her degree of confidence by the Bayesian rule of conditionalization. Once we adopt this Bayesian framework, we expect the aggregate degrees of confidence the proportional weight view recommends will also satisfy the Bayesian constraints. It turns out they don't, as we will see now.

Since the proportional weight view leads to consensus degrees of confidence,[12] I will write the aggregate degree of confidence in *p* as $C^*(p)$, dropping the subscript for the

---

[9] It is assumed throughout the paper that the disagreeing people make their individual judgments independently of each other before they take into account of other people's judgments.

[10] See Lehrer and Wager (1981) for a full account of this approach.

[11] One nice feature of the proportional weight view is that it leads to a consensus—i.e. $C_1^*(p) = \ldots = C_n^*(p) = q_1C_1(p) + \ldots + q_nC_n(p)$—as long as the disagreeing people agree on their epistemic qualifications. This is helpful when the disagreeing people need to make a joint decision in a cooperative venture (Gillies 1991, 2000 Ch. 8; Gillies and Ietto-Gillies 1991). Meanwhile, those who give a bigger weight to their own judgment than their qualification warrants fail to achieve a consensus even if they can agree on their epistemic qualifications. However, this may not be a big advantage after all, at least theoretically. DeGroot (1974) and Lehrer (1975) present a version of the Delphi Method in which *repeated* applications of the weighted averaging of different degrees of confidence converge under certain general conditions even if they disagree on each other's epistemic qualifications.

[12] See note 11 above.

epistemic subject. The Bayesian constraints on the aggregate degrees of confidence are then: C* is a probability function and updated in accordance with the Bayesian rule of conditionalization, i.e. $C^{+r}*(p) = C*(p|r) =_{def.} C*(p \& r)/ C*(r)$. The following example shows that the version of the proportional weight view considered in Section 1 violates the latter.

Example 1: Let $C_1(h \& e) = 1/6$, $C_1(e) = 1/3$, $C_2(h \& e) = 1/6$, $C_2(e) = 1/2$, and Q = <.5, .5>. It follows from these assignments that:

$$C_1^{+e}(h) = C_1(h|e) = C_1(h \& e)/C_1(e) = 1/2$$

$$C_2^{+e}(h) = C_2(h|e) = C_2(h \& e)/C_2(e) = 1/3$$

$$\therefore C^{+e}*(h) = (C_1^{+e}(h) + C_2^{+e}(h))/2 = 5/12$$

Meanwhile,

$$C*(h \& e) = (1/6 + 1/6)/2 = 1/6$$

$$C*(h) = (1/3 + 1/2)/2 = 5/12$$

$$\therefore C*(h|e) = C*(h \& e)/C*(h) = (1/6)/(5/12) = 2/5$$

Hence, $C^{+e}*(h) \neq C*(h|e)$ in violation of the Bayesian rule of conditionalization.[13]

[13] C* also fails to preserve probabilistic independence, i.e. even if $C_i(p|r) = C_i(p)$ for all $i = 1, \ldots, n$, it may still not be the case that $C*(p|r) = C*(p)$. To see this, add to the example above, $C_1(h) = 1/2$ and $C_2(h) = 1/3$. Under these assignments $C_1(h) = C_1(h|e) = 1/2$ and $C_2(h) = C_2(h|e) = 1/3$, but $C*(h) = 5/12$ while $C*(h|e) = 2/5$. This means that even though S_1 and S_2 agree that e is irrelevant to h, e disconfirms (lowers the probability of) h according to their aggregate degrees of confidence.

This is a problem, but it is only a counterexample to the proportional weight view *as formulated in Section 1*, which takes the mean to be the arithmetic mean. This is not the only option for the proportional weight view since there are many other means known to mathematicians. For example, the three classic "Pythagorean" means are the arithmetic mean (A), the geometric mean (G), and the harmonic mean (H):

$$A(x_1, \ldots, x_n) = (\sum_{i=1}^{n} x_i)/n$$

$$G(x_1, \ldots, x_n) = (\prod_{i=1}^{n} x_i)^{1/n}$$

$$H(x_1, \ldots, x_n) = n/(\sum_{i=1}^{n} 1/x_i)$$

It turns out that the proportional weight view obeys the Bayesian rule if we adopt the weighted geometric mean G_Q instead of the weighted arithmetic mean A_Q, as shown below, where the weights $Q = \langle q_1, \ldots, q_n \rangle$ are given in the form of exponents:

$$\begin{aligned}
C^{+r}*(p) &= G_Q(C_1^{+r}(p), \ldots, C_n^{+r}(p)) \\
&= \prod^n{}_{i=1} C_i^{+r}(p)^{q_i} \\
&= \prod^n{}_{i=1} C_i(p|r)^{q_i} \\
&= \prod^n{}_{i=1} [C_i(p \,\&\, r)/C_i(r)]^{q_i} \\
&= \prod^n{}_{i=1} [C_i(p \,\&\, r)^{q_i}/C_i(r)^{q_i}] \\
&= \prod^n{}_{i=1} C_i(p \,\&\, r)^{q_i}/\prod^n{}_{i=1} C_i(r)^{q_i} \\
&= G_Q(C_1(p \,\&\, r), \ldots, C_n(p \,\&\, r))/G_Q(C_1(r), \ldots, C_n(r)))
\end{aligned}$$

$= C^*(p \,\&\, r)/C^*(r)$

However, abandoning the weighted arithmetic mean in favor of the weighted geometric mean does not save the proportional weight view because C* as determined by the weighted geometric mean is not a probability function even if C_i's are. In order to be a probabilistic function, C* must satisfy the following three constraints:

Non-Negativity: $0 \leq C^*(p)$ for all p.

Normalization: If p is a tautology, $C^*(p) = 1$.

Additivity: If p and r are logically incompatible, $C^*(p \vee r) = C^*(p) + C^*(r)$

The following example shows that C* as determined by the weighted geometric mean fails to satisfy the constraints.

Example 2: Let $C_1(e) = 1/3$, $C_2(e) = 1/2$, and Q = <.5, .5>. It follows from these assignments that:

$$
\begin{aligned}
C^*(e) &= G_Q(C_1(e), C_2(e)) \\
&= [C_1(e) \times C_2(e)]^{1/2} \\
&= [1/3 \times 1/2]^{1/2} \\
&= (1/6)^{1/2}
\end{aligned}
$$

$$C^*(\neg e) = G_Q(C_1(\neg e), C_2(\neg e))$$

$= [C_1(\neg e) \times C_2(\neg e)]^{1/2}$

$= [(1 - C_1(e)) \times (1 - C_2(e))]^{1/2}$

$= [(1 - 1/3) \times (1 - 1/2)]^{1/2}$

$= [2/3 \times 1/2)]^{1/2}$

$= (1/3)^{1/2}$

Meanwhile, $C^*(e \vee \neg e) = 1$ since $e \vee \neg e$ is a tautology. So, $C^*(e \vee \neg e) \neq C^*(e) + C^*(\neg e)$ in violation of the additivity constraint.[14]

It turns out that the original formulation of the proportional weight view that uses the weighted arithmetic mean A_Q does well in this regard. It is trivial that the credence function obtained by the weighted arithmetic mean satisfies of the non-negativity constraint, i.e. $0 \leq C^*(p)$ as long as $0 \leq Ci(p)$ for all $i = 1, \ldots, n$. It is also trivial that it satisfies the normalization constraint, i.e. $C^*(p) = 1$ as long as $C_i(p) = 1$ for all $i = 1, \ldots, n$. The only serious question is whether C* obtained by the weighted arithmetic mean satisfies the additivity constraint, i.e., $C^*(p \vee r) = C^*(p) + C^*(r)$ on condition that $C_i(p \vee r) = C_i(p) + C_i(r)$ for all $i = 1, \ldots, n$, and we can see this easily as follows. (In the case of the weighted arithmetic mean, the weights $Q = \langle q_1, \ldots, q_n \rangle$ are given in the form of multipliers.)

[14] A related problem about the (weighted) geometric mean is that if one person assigns zero to e, then the aggregate degree of confidence in e is automatically zero regardless of other people's assignments. If another person assigns zero to ¬e, then the aggregate degree of confidence in ¬e is also zero. This means that C*(e) + C*(¬e) becomes zero instead of one, as it should be.

$$
\begin{aligned}
C^*(p \vee r) &= A_Q(C_1(p \vee r), \ldots, C_n(p \vee r)) \\
&= \textstyle\sum^n_{i=1} [q_i C_i(p \vee r)] \\
&= \textstyle\sum^n_{i=1} [q_i(C_i(p) + C_i(r))] \\
&= \textstyle\sum^n_{i=1} [q_i C_i(p) + q_i C_i(r)] \\
&= \textstyle\sum^n_{i=1} [q_i C_i(p)] + \sum^n_{i=1} [q_i C_i(r)] \\
&= A_Q(C_1(p), \ldots, C_n(p)) + A_Q(C_1(r), \ldots, C_n(r)) \\
&= C^*(p) \vee C^*(r)
\end{aligned}
$$

Here is what we uncovered so far. If we adopt the weighted arithmetic mean to implement the proportional weight view, then C* is a probability function but C* violates the Bayesian rule of conditionalization. On the other hand, if we adopt the weighted geometric mean to implement the proportional weight view, then C* obeys the Bayesian rule of conditionalization but C* is not a probability function. So, neither of them is satisfactory from the Bayesian standpoint. It is still too early for the Bayesian to give up the proportional weight view since there are many other means that we haven't examined, such as the harmonic mean and the quadratic mean (also known as the root mean square). There may be an appropriate mean, M, which makes the proportional weight view works, i.e. C* obtained by the weighted M mean is a probability function and obeys the Bayesian rule of conditionalization.

However, it became known in the early eighties that under a very plausible assumption, no mean—and more generally no systematic aggregation of individual degrees of confidence—satisfies the two Bayesian constraints.[15] The plausible

[15] See Genest and Zidek (1986) for a review of the literature.

assumption, the Context-Free Assumption, is that the aggregate degree of confidence on any proposition *p* is a function of the individual degrees of confidence on that proposition *p*, i.e., $C^*(p) = F(C_1(p), \ldots, C_n(p))$ for some function F.[16] This means that we can determine the aggregate degree of confidence on a proposition *locally* without consulting the individual degrees of confidence on other propositions. McConway (1981) and Wagner (1982) showed independently that under the Context-Free Assumption *only* C* as determined by the weighted arithmetic mean is a probability function.[17] Since C* as determined by the weighted arithmetic mean violates the Bayesian rule of conditionalization, it is impossible, under the Context-Free Assumption, to satisfy the two Bayesian constraints. I call it the Bayesian double bind.

3. Options and Further Issues

Some opponents of the proportional weight view may take the Bayesian double bind to be a vindication of their position, i.e., incompatibility with Bayesianism is a good reason to reject the proportional weight view. However, that depends on what alternative view they support. If they take the position of "My way or the highway" in total disregard for other people's judgments, they can embrace the Bayesian double bind, and there may be some support for this position as an alternative to the (allegedly spineless) equal weight

---

[16] It is also assumed here that the function F is not dependent on the proposition *p*. In other words, it is not the case that different functions are used for different propositions.

[17] To be more precise, the proof assumes that the aggregate degrees of confidence for three or more propositions are determined. In other words, the case involves more than one proposition and its negation.

view. However, it is an unreasonable position as an alternative to the proportional weight view, for it means that even a total novice should make no adjustment at all in her degree of confidence when all experts (whom she acknowledges to be her epistemic superiors) disagree with her. A more reasonable alternative is the *disproportional* weight view that gives a higher weight to one's own judgment than one's epistemic qualification warrants without completely ignoring other people's judgments. However, the Bayesian double bind is just as troubling for this moderate alternative because the double bind applies to any systematic weighting of disagreeing judgments, whether the weights are proportional or disproportional to the epistemic qualifications. We can see this by simply replacing the epistemic qualifications $Q = \langle q_1, \ldots, q_n \rangle$ with the disproportional weights $W = \langle w_1, \ldots, w_n \rangle$ throughout the discussion of the Bayesian double bind. We still get the same results, i.e. skewed weighting in favor of one's own judgment is still incompatible with Bayesianism.

It is not quite accurate to say that total disregard for other people's judgment is the only view consistent the Bayesian double bind. One can instead totally ignore one's own judgment in complete deference to someone else's judgment, but one has to defer to the judgment of a single person.[18] In other words, the Bayesian double bind forces us to pick a single person, either oneself or someone else, and ignore all others. The attitude is "My way or her way" with no mutual accommodation and in total disregard for all other judgments. Some may think it fine to pick the best-qualified expert and follow her way, but when the equally qualified experts in the field disagree among themselves, why does one have to pick a single expert and ignore others? It is more sensible to seek a balanced

[18] One can defer to the judgment of a group of people who are in complete agreement among themselves, but that is no different from complete deference to the judgment of a single person in the group.

view by giving appropriate weights to their different opinions. If it so happens that I am among the best experts, holding on to my own view may make some sense in consideration of the spineless concern, but when I am an epistemic inferior seeking expert opinions, there is really no good reason to single out one expert and ignore all others.

Given the unpalatable choice of "My way or her way" it is worthwhile to revisit the assumptions that put us in this predicament. An obvious assumption to question is Bayesianism itself, but the Context-Free Assumption may also be reconsidered. Recall that the Context-Free Assumption states that there is a function F such that:

$$C^*(p) = F(C_1(p), \ldots, C_n(p))$$

It turns out that there is a natural way of relaxing the assumption that allows us to avoid the Bayesian double bind. The relaxed assumption is:

$$C^*(p) \propto F(C_1(p), \ldots, C_n(p)),$$

where $\propto$ means proportionality. More specifically, the relaxation allows the normalization of the weighted geometric mean by the sum of all the weighted geometric means. Here is the way it works. Let $p_1, \ldots, p_m$ be pairwise incompatible, jointly exhaustive propositions. The weighted geometric mean of the individual degrees of confidence for p_k is $\prod^n_{i=1} C_i(p_k)^{q_i}$. So, the sum of all the weighted geometric means for p_1,

…, p_m is $\sum^m_{j=1}(\prod^n_{i=1}C_i(p_j)^{q_i})$. If we use this sum to normalize the weighted geometric mean for each proposition, the aggregate degree of confidence $C^*(p_k)$ is:

$$C^*(p_k) = \prod^n_{i=1}C_i(p_k)^{q_i} / \sum^m_{j=1}(\prod^n_{i=1}C_i(p_j)^{q_i})$$

Thanks to normalization, C* satisfies the additivity constraint; and being proportional to the weighted geometric mean, C* also satisfies the Bayesian rule of conditionalization.[19]

It may seem normalization does the magic and we can now get out of the Bayesian double bind, but there are some serious problems about the normalization strategy. First, it leads to the violation the Unanimity Principle, i.e., it is no longer the case that if all individuals share the same degree of confidence for proposition p_k, then that is the aggregate degree of confidence for p_k. Second, it makes the procedures of aggregation and marginalization non-commutative.[20] In other words, aggregated marginals and marginalized aggregates are not necessarily the same. We can see both of these problems in the following example.

Example 3: Let S_1 and S_2 be epistemic peers (i.e. $q_1 = q_2 = .5$), and let C_1 and C_2 be such that:

$C_1(\neg a \;\&\; b) = C_1(\neg a \;\&\; \neg b) = C_2(\neg a \;\&\; b)$, $C_2(\neg a \;\&\; \neg b) = .25$

$C_1(a \;\&\; b) = .2$, $C_2(a \;\&\; \neg b) = .3$

[19] In fact Genest (1984) shows that this is the only way to satisfy the two Bayesian constraints under the relaxed Context-Free Assumption.

[20] See McConway (1981).

$C_2(a \& b) = .3$, $C_2(a \& \neg b) = .2$

To begin with the violation of the Unanimity Principle, note that the two weighted geometric means $G_Q(C_1(\neg a \& b), C_2(\neg a \& b))$ and $G_Q(C_1(\neg a \& \neg b), C_2(\neg a \& \neg b))$ are simply .25 before normalization because of the agreement (unanimity) between S_1 and S_2 on these propositions. Meanwhile $G_Q(C_1(a \& b), C_2(a \& b))$ and $G_Q(C_1(a \& \neg b), C_2(a \& \neg b))$ are less than .25 because of the disagreement between S_1 and S_2 on these propositions. As a result, the sum of all the weighed geometric means is less than 1. This means that the normalization by the sum of all the weighed geometric means makes $C^*(\neg a \& b)$ and $C^*(\neg a \& \neg b)$ greater than .25 despite the unanimity $C_1(\neg a \& b) = C_2(\neg a \& b)) = .25$ and $C_1(\neg a \& \neg b) = C_2(\neg a \& \neg b)) = .25$. The Principle of Unanimity is violated.[21]

The example also shows that normalization makes aggregation and marginalization non-commutative. Note first that $C_1(a) = C_2(a) = .5$ despite $C_1(a \& b) \neq C_2(a \& b)$ and $C_1(a \& \neg b) \neq C_2(a \& \neg b)$. There are two ways of calculating $C^*(a)$. One way is to calculate $C^*(a \& b) < .25$ and $C^*(a \& \neg b) < .25$ by normalizing the weighted geometric means; and then marginalize them to obtain $C^*(a) < .5$. The other way is to calculate $C_1(a) = .5$ and $C_2(a) = .5$ by marginalizing, respectively, $C_1(a \& b)$ and $C_1(a \& \neg b)$, and $C_2(a \& b)$ and $C_2(a \& \neg b)$; and then aggregate them to obtain $C^*(a) = .5$. The two results—the marginalized aggregates and the aggregated marginals—are different. This is troubling because it means that we obtain different aggregate degrees of

[21] The problem in a more general form is that $C^*(p)$ can be lager than the largest among $C_1(p), \ldots, C_n(p)$, or smaller than the smallest among $C_1(p), \ldots, C_n(p)$. The violation of the Unanimity Principle is a special case where $C_1(p) = \ldots = C_n(p)$.

confidence for the same proposition, depending on when we come to focus our attention on the proposition—before aggregation or after aggregation.[22]

Normalization is not an easy way out after all, and we need to reconsider the Bayesian rule of conditionalization. Note first that we need not question the Bayesian rule of conditionalization in general, but only the Bayesian rule of conditionalization as applied to the aggregate degrees of confidence. In fact one way to understand C*'s violation of the Bayesian rule of conditionalization is that the weighted arithmetic mean (the only aggregation allowed by the Context-Free Assumption) makes aggregation and Bayesian updating non-commutative. When we aggregate individual degrees of confidence by the weighted arithmetic means, and then update them by the Bayesian rule of conditionalization, we get one result. When we update individual degrees of confidence by the Bayesian rule of conditionalization, and then aggregate them by the weighted arithmetic means, we can get a different result. This means that we have two choices—i.e. take either the updated aggregates or the aggregated updates, but not both.

The first choice means that we only update aggregates, and never aggregate updated individual degrees of confidence. The aggregate posterior degree of confidence for *p* would be simply the ratio, $C^*(p \,\&\, r)/C^*(r)$, and not the aggregation of individual posterior degrees of confidence $C_i^{+r}(p)$. The problem with this policy is that it makes the result dependent on the timing of aggregation. Suppose, for example, one person S_1 was aware of the difference in people's degrees of confidence before the truth of e is known,

[22] Marginal unanimity in the example, $C_1(a) = C_2(a)$, is not essential for the failure of the commutativity between aggregation and marginalization. Even if $C_1(a)$ and $C_2(a)$ are different, the degree to which the weighted geometric mean makes C*(a) different from their weighted arithmetic mean, are not generally the same as the degree to which the weighted geometric mean makes C*(a & b) and C*(a & ¬b) different from their weighted arithmetic means, and the difference can remain even after normalization. So, C*(a) obtained from $C_1(a)$ and $C_2(a)$ can still be different from the sum of C*(a & b) and C*(a & ¬b).

while another person S_2 notices the difference only after the truth of e is known. S_1 would have calculated C*(h & e) and C*(e) before the truth of e is known, and thus calculate the ratio, C*(h & e)/C*(e), to obtain the aggregate posterior degree of confidence once the truth of e is known, according to the proposed policy. For S_2, on the other hand, the aggregate posterior degree of confidence is the weighted arithmetic mean of the individual posterior degrees of confidence since her aggregation takes place after the truth of e is known. S_2 only notices other people's updated degrees of confidence, C_1^{+e}(h), C_3^{+e}(h), etc., but not their prior degrees of confidence.[23] So S_2 cannot calculate C*(h & e) and C*(e) to obtain C*(h & e)/C*(e), as S_1 does. The resulting aggregate degrees of confidence can be different, then, depending on the timing of aggregation, which is very undesirable.

The second choice, on the other hand, means that each time a new piece of evidence becomes available we consult updated individual degrees of confidence and aggregate them. We do not update aggregate degrees of confidence, so $C^{+r}*(p)$ is not the same as the aggregate conditional probability C*(*p*|*r*), or C*(*p* & *r*)/C*(*r*). One obvious problem is that this makes us vulnerable to a diachronic Dutch Book.[24] However, the assessment is a matter of comparison at this point and if other alternatives look even worse, we might as well pay this price. One could also argue that we *should* ignore C*(*p*|*r*) in updating aggregate degrees of confidence because C*(*p*|*r*) violates the Unanimity Principle as applied to conditionals, i.e. C*(p|r) can be different from a

[23] S_2 should be able to tell C_1(a & b)/C_1(b) since it should be equal to C_1^{+}(a) if S_1 obeys the Bayesian rule of conditionalization, but S_2 cannot tell C_1(a & b) and C_1(b) from C_1(a & b)/C_1(b).

[24] See Teller (1973).

unanimous conditional probability, $C_1(p|r) = \ldots = C_n(p|r)$.[25] In other words, we have reason to question C*(p|r) even before we update our degrees of confidence. However, the observation that $C^*(p|r)$ violates the Unanimity Principle as applied to conditionals raises the question whether we are addressing the problem at the right place. The observation reveals that the problem does not arise suddenly when we update the aggregate degrees of confidence. Even before we obtain new evidence and update the aggregate degrees of confidence, the Unanimity Principle is violated with regard to conditionals. Seen in this way, the real problem is not updating the aggregate degrees of confidence but the aggregate conditional degrees of confidence. Problem, in other words, is that when we aggregate the individual degrees of confidence by the weighted arithmetic means (the only aggregation allowed under the Context-Free Assumption), we cannot respect the Unanimity Principle for conditional degrees of confidence.

The question we need to ask now is whether we need to respect the Unanimity Principle with regard to conditional degrees of confidence, and the answer seems to be different in different cases.[26] In some cases, we want to make the unanimous conditional degree of confidence to be the aggregate conditional degree of confidence. Suppose we all have good reason to assign the same conditional probability $C_1(p|r) = \ldots = C_n(p|r)$ because of the relation between *p* and *r*. We then want $C^*(p|r)$ to be $C_1(p|r) = \ldots = C_n(p|r)$, respecting unanimity. But in other cases unanimity may be an accident. Suppose S_1 and S_2 assigns different degrees of confidence $C_1(p) \neq C_2(p)$ and $C_1(p \vee r) \neq C_2(p \vee r)$, but the ratio $C_1(p)/C_1(p \vee r)$ happens to be the same as the ratio $C_2(p)/C_2(p \vee r)$. This

[25] See Dalkey (1972, 1975).

[26] For analogous discussions on the question of whether aggregation should preserve unanimity on *probabilistic independence*, see Laddaga (1977), Lehrer and Wagner (1983), Loewer and Laddaga (1985), and Wagner (1985).

means that $C_1(p|p \vee r) = C_2(p|p \vee r)$ since *p* is logically equivalent to *p* & ($p \vee r$). In this case there happens to be unanimity on the conditional degrees of confidence, but I don't think it is necessary to make $C^*(p|p \vee r)$ to be the same as $C_1(p|p \vee r) = C_2(p|p \vee r)$.

One way to generalize the difference is that some conditional degrees of confidence are "basic" in the sense that they do not depend on other degrees of confidence, while other conditional degrees of confidence are "derived" from other degrees of confidence in consideration of the Bayesian constraints. Once we make this distinction it seems reasonable to respect unanimity in conditional degrees of confidence when they are basic, but not when they are derived from other degrees of confidence. The distinction between basic and derived degrees of confidence may also apply to marginal and joint degrees of confidence. We may respect unanimity in *any* degrees of confidence if they are basic, but not if they are derived from other degrees of confidence. This points to a general response to the Bayesian double bind beyond the preservation of the unanimity, viz. we only aggregate basic degrees of confidence, and not derived degrees of confidence. This allows us to avoid violating the Bayesian constraints in many troubling cases.

Needless to say, there remain many issues to address even if this is the right way of responding to the Bayesian double bind. It is somewhat unclear, for example, why we should refrain from aggregating derived degrees of confidence. The idea has some intuitive appeal, but the main motivation is to avert troubles. It may be considered a forced maneuver. There are also more specific problems to deal with. For example, the trouble remains in cases where all degrees of confidence are basic, including conditional degrees of confidence, though this may not be a realistic possibility. Also, there may not be a clear distinction between basic and derived degrees of confidence in some cases. Our degrees of confidence often result from some sort of reflective equilibrium, or mutual adjustment among conflicting degrees of confidence, in which case no resulting degrees of confidence are completely basic or completely derived. There may be a sensible way of resolving these issues, but it is also quite possible that the Bayesian double bind is an indication that epistemology of disagreement, which is still in an early stage of development, overlooks something important, and that we may need a radically different approach.

Reference

Aumann, R. 1976. “Agreeing to Disagree” *The Annals of Statistics* 4 (6), pp. 1236-1239.

Christensen, D. 2007. “Epistemology of Disagreement” *Philosophical Review* 116 (2), pp. 187-217.

Dalkey, N. 1972. “An Impossibility Theorem for Group Probability Functions” Santa Monica CA: RAND Corporation, P-4862.

Dalkey, N. 1975. “Toward a Theory of Group Estimation” in H. A. Linstone and M. Turoff (eds.) *The Delphi Method: Techniques and Applications* (Reading, MA: Addison-Wesley), pp. 231-256.

DeGroot, M. (1974) “Reaching a Consensus” *Journal of the American Statistical Association* 69, pp. 118-121.

Elga, A. Forthcoming. “Reflection and Disagreement” *Noûs*.

Genest C. 1984. “A Characterization Theorem for Externally Bayesian Groups” *Annals of Statistics* 12, pp. 1100-1105.

Genest, C. and Zidek, J. 1986. “Combining Probability Distributions: A Critique and an Annotated Bibliography,” *Statistical Science* 1, pp. 114-135.

Gillies, D. 1991. “Intersubjective Probability and Confirmation Theory,” *British Journal for the Philosophy of Science* 42 (4), pp. 513-533.

Gillies, D. and Ietto-Gillies, G. (1991), “Intersubjective Probability and Economics,” *Review of Political Economy* 3 (4), pp. 393-417.

Gillies, D. 2000. *Philosophical Theories of Probability* (London: Routledge).

Kelly, T. 2005. “The Epistemic Significance of Disagreement” in J. Hawthorne and T. Gendler (eds.) *Oxford Studies in Epistemology* Vol. 1.

Kitcher, P. 1990. “The Division of Cognitive Labor,” *The Journal of Philosophy*, 87 (1), pp. 5-22.

Laddaga, R. 1977. “Lehrer and the Consensus Proposal,” *Synthese* 36, pp. 473-477.

Lehrer, K. 1975. “Social Consensus and Rational Agnoiology,” *Synthese* 31, pp. 141-160.

Lehrer, K. and Wagner, C. 1981. *Rational Consensus in Science and Society: A Philosophical and Mathematical Study* (Dordrecht, Holland: Reidel Publishing).

Lehrer, K. and Wagner, C. 1983. “Probability Amalgamation and the Independence Issue: A Reply to Laddaga” *Synthese* 55, pp. 339-346.

Levi, I. 1980. *The Enterprise of Knowledge: An Essay of Knowledge, Credal Probability, and Chance* (Cambridge MA: MIT Press).

Loewer, B. and Laddaga, R. 1985. “Destroying the Consensus,” *Synthese* 62, pp. 79-95.

McConway, K. J. 1981. “Marginalization and Linear Opinion Pools,” *Journal of American Statistics* 76, pp. 410-414.

Teller, P. 1973. “Conditionalization and Observation,” *Synthese* 26 (2), pp. 218-258.

van Inwagen, P. 1996. “It Is Wrong, Everywhere, Always, and for Anyone, to Believe Anything upon Insufficient Evidence” in Jeff Jordan and Daniel Howard-Snyder (eds.) *Faith, Freedom, and Rationality: Philosophy of Religion Today* (London: Rowman and Littlefield).

Wagner, C. 1982. “Allocation, Lehrer Models, and the Consensus of Probabilities,” *Theory and Decision* 14, pp. 207-220.

Wagner, C. 1985. “On the Formal Properties of Weighted Averaging as a Method of Aggregation,” *Synthese* 25, pp. 233-240.