Theory and Decision manuscript No.
(will be inserted by the editor)

Learning from Others

Conditioning versus Averaging

Richard Bradley

Department of Philosophy, Logic and Scientific Method
London School of Economics, Houghton Street, London WC2A 2AE

The date of receipt and acceptance will be inserted by the editor

Abstract In the face of disagreement in the expressed probabilities of one or more individuals on some proposition, it has been suggested that we should revise our beliefs by adopting a linear average of the expressed opinions on it. Is such belief revision compatible with Bayesian conditionalisation? In this paper I look at situations in which full or partial deference to the expressed opinions of others is warranted to consider what Bayesianism and linear averaging respectively require of us. I will conclude that only in trivial circumstances are the requirements imposed by the two compatible.

Key words Linear averaging, Bayesian conditionalisation, expert testimony, deference

1 Introduction

When others have information or judgemental capabilities that we lack, then their opinions are a resource that we can and should exploit for the purposes of forming or revising our own opinions. But how should we do this? In this paper I will compare two types of answer to this question – Bayesian conditioning and opinion pooling – and ask whether they are compatible. Some interesting work on the question suggests a positive answer under various conditions: See for instance Genest and Schervish [10], Bonnay and Cozic [3][4] and Romeijn [17]. But how restrictive these conditions are remains an open question.

A Bayesian treats the expressed opinions of others as evidence for and against the truth of the claim under consideration, evidence whose relevance is captured by an assignment of a conditional probability for the claim, given each possible combination of others' opinions. She responds to this evidence by conditioning on it, i.e. by adopting as her revised opinion her (prior) conditional degrees given the expressed opinions. The opinion pooler, on the other hand, adopts as her new opinion an aggregate of the expressed opinions of others, an aggregate that in some way reflects the epistemic value that she attaches to each of the expressed opinions.

I shall assume here that Bayesianism provides the gold standard for coherent revision of belief in the kinds of situation in which it applies, namely when we have prior probabilities for not only the hypotheses of ultimate interest but also for all possible combinations of evidence (in this case, the expressions of opinion) that either confirm or disconfirm these hypotheses, and when everything that we learn is representable by one such possible combination. The problem is that it is not always easy to apply the Bayesian theory. In circumstances in which the evidence takes a 'non-standard' form, such as when it is imprecise or conditional in form, we can turn to other forms of belief revision, such as Jeffrey conditioning or Adams conditioning.[1] But when we are unable to assign prior probabilities to the possible evidence propositions then no conditioning method at all may be applicable. This can happen because there are simply too many possibilities for us to process them all, or because we don't have enough information to assign a precise probability with any confidence.

These difficulties are acutely pertinent to the question of how to exploit the information taking the form of expert opinion reports. Consider, for instance, someone who claims to be an expert on wines. What's the probability that they will make any particular judgement about any particular wine? If you don't know much about wine, it will be hard to say. In the statistics literature, agents who change their beliefs by conditionalising on the testimonial evidence of experts are known as supra-Bayesians (see, for instance, [16] and [9]). Supra-Bayesians must have priors over the opinion states of all those whose opinions count as evidence for them with regard to some proposition. But opinions about opinions might be evidence too, and opinions about opinions about opinions. And so on. It would be fair to say that supra-Bayesians are required to be cognitive super-Humans.

A Bayesian with more limited cognitive resources has two reasons for taking an interest in opinion pooling. Firstly, it might help her in thinking about how to assign the probabilities to the hypotheses, conditional on combinations of opinion, that she needs in order to revise her own opinions by conditionalisation when she gets information of this kind. Secondly, in circumstances in which she cannot conditionalise on expressions of opinion because she lacks the requisite conditional probabilities for the hypotheses that interest her, she might adopt opinion pooling as an alternative method of belief revision. In both cases, the Bayesian will want to know whether a rule for opinion pooling is compatible with her commitment to conditionalisation. This is true not just when she uses opinion pooling as a guide to making conditional probability judgements, but also in the case when she

[1] See [5] for a discussion.

uses it as an alternative to conditionalisation. For in this latter case, she will want to know that there is some way of assigning prior probabilities to combinations of expressed opinion and to the hypotheses that are evidentially dependent on them, such that the pooled opinion is what she would have obtained by conditionalisation if these had been her prior probabilities.

In this paper I will address the question of whether revision of belief in response to testimonial evidence by application of a specific form of opinion pooling – linear averaging – is compatible with Bayesian norms of belief change. I will begin by making more precise what is required for opinion pooling rules to be compatible with Bayesian revision and rehearse some old arguments for why unrestricted linear averaging fails to meet the compatibility criterion. In subsequent sections, I consider linear averaging on restricted domains, drawing particularly on Dawid et al [6], to test for compatibility with Bayesianism. I do so by exploring what Bayesianism requires of agents in some rather simple situations; in particular ones which license deference to the opinions of one or more experts on some specific proposition on grounds of their superior knowledge. The conclusion is somewhat surprising. Linear averaging, even when applied to a single proposition, is not (non-trivially) compatible with Bayesian norms in situations involving more than one source of testimony.

2 Linear Averaging and Bayes Compatibility

To formulate our questions in a general way, let $\mathcal{S}$ be a Boolean algebra of propositions containing at least four propositions and Ω be a finite subset of $\mathcal{S}$ containing at least one non-contradictory proposition, serving here to represent the set of propositions regarding which our protagonist—the Bayesian agent—will learn the opinions of others. Let an opinion function f on Ω be a mapping from Ω to $[0, 1]$ that is extendible to a probability function on Ω. Let Π be the set of all such opinion functions on Ω. Then a mapping $\Phi : \Pi^n \to \Pi$, from a profile of n opinion functions on Ω to an 'aggregate' opinion function, is called an opinion formation rule for Ω. Intuitively such a rule tells the agent what opinions to adopt on the propositions in Ω when she learns those of the n others.

In this paper we will look only at one particular class of opinion formation rules: the linear averaging rules. A *linear averaging rule* for Ω is a mapping F from a profile of opinion functions $(f_1, ..., f_n)$ on the propositions in Ω to an average opinion for them such that for some set of non-negative weights $\{\alpha_i\}$ such that $\sum_i \alpha_i = 1$, for all $X \in \Omega$:

$$F(f_1, ..., f_n)(X) = \sum_i \alpha_i . f_i(X)$$

There is a large statistics literature, and a modest philosophical one, on linear averaging and its properties, most of it focused on the case in which opinion functions are defined on full Boolean algebras and so are forcibly probability functions: see Genest and Zidek [11] and Dietrich and List [7] for good surveys. In this context, the weights α_i are canonically interpreted as measures of the (relative) judgemental competence of individuals. And linear pooling is thought of as a way of constructing an aggregate probability for each proposition X by taking a competence-weighted average of the probabilities assigned to it by each individual. Much of this literature is focused on the problem of how to form a *consensual* opinion on some issue, a question that is not directly related to that of how individuals should improve their own opinions by taking into account the information provided by the opinions of others. These two issues can come together, as they do in the theory of Lehrer and Wagner (see [19], [15] and [20]), if the way that the consensus is achieved is by individuals revising their beliefs rationally in response to the expressed opinions of others. But here we will concern ourselves only with the question of whether revising one's beliefs in this way is epistemically rational and set aside the issue of whether or not it will lead to a consensus.

This is important because whatever the merits of linear averaging as a means of forming a consensual or group probability from a diverse set of individual ones, it is seriously flawed as a method of belief revision (contrary to Lehrer's [14] claim that rationality requires that we revise our beliefs in this way). This is so for two important reasons. Firstly, linear averaging is not sensitive to the proposition-specific competencies of individuals. But intuitively individuals have different domains of expertise: I don't take my plumber's pronouncements about health risks very seriously, for instance, but equally I would not get my doctor to tell me the cause of the drainage problems in my house. It follows that any rule that uses constant competence weights will be unsatisfactory to some degree.

Secondly, linear averaging is insensitive to whether the opinions expressed by different individuals on the same proposition are independent or not. But compare a situation in which two scientists conduct separate experiments to try and settle some question with one in which they conduct a single experiment together. Suppose that in both cases the scientists report that as a result of their experiments they consider X to be highly probable. In the former case, we would want to raise our own probability for X quite considerably because of the convergence of their expert testimony. In the latter case too we would want to raise our probability for X, but less so, because their joint testimony in favour of X is based on same information. To revise once in the light of the testimony of the first scientist and then again in the light of that of the second would in effect be to update twice on the same evidence, akin to an individual scientist conditioning twice on the same experimental result.

In contrast, a supra-Bayesian is free to assign a conditional probability to a proposition, given the testimony of some expert, that reflects her judgement of the competence of that expert on that proposition. Furthermore, her conditional probabilities given the joint testimonies of the experts will typically depend on the probabilistic dependencies between expert reports.

It is not surprisingly therefore that linear averaging on an unrestricted domain is compatible with Bayesianism in only very special circumstances: In particular, when the opinions of all but one of the experts consulted are independent of the truth (for a proof of this claim see [2]), a circumstance that rather undermines the point of consulting them in the first place.

Now, as Steele [18] points out, these problems might not arise when linear averaging is applied to restricted sets of propositions—in particular, partitions—for in this case we are free to apply domain-dependent weights that reflect both the specific competence of the individuals on that domain and any dependencies between their opinions within it. For instance, if three individuals express an opinion on the same proposition, but two of them have formed it in discussion with each other, then we can choose to attach a large weight to the pair of them, because their deliberations increase their competence, but count them as a single unit in the linear average.

Steele's suggested restriction is insufficiently severe, however, because linear averaging on a partition implies constant weights over any coarsening of it. Consider the following example. Suppose that You know that Anne has observed that A is true while Bob has observed that B is true. Suppose furthermore that they report the following probabilistic degrees of belief across the partition $\pi = \{AB, A\neg B, \neg AB, \neg A\neg B\}$:

	AB	$A\neg B$	$\neg AB$	$\neg A\neg B$
Anne	0.1	0.9	0	0
Bob	0.1	0	0.9	0
Linear average	0.1	$0.9a$	$0.9b$	0

Now Anne and Bob's observations make them maximally competent, respectively, on the question of whether or not A is true and whether or not B is true (supposing absence of observational error). So You should simply adopt their reported beliefs as your own, leaving You with degree of belief 1 in AB and 0 in all the other propositions. But this goes against the recommendations of linear averaging which, irrespective of the weights a and b assigned to Anne and Bob, will yield a probability of 0.1 for the proposition AB and 0 for $\neg A\neg B$.

We can put the problem slightly differently. If You revise your beliefs by linear averaging in response to Ann and Bob's testimony on the four element partition π using weights a and b for Anne and Bob respectively, then your new degrees of belief on the partitions $\{A, \neg A\}$ and $\{B, \neg B\}$ must, on pain of probabilistic incoherence, be linear averages of Anne's and Bob's degrees of belief obtained by applying weights a and b. But Anne and Bob have different competencies over these partitions, so these weights cannot be adequate to both. The upshot is that if linear averaging is to serve the purposes of rational belief revision in response to testimony, then it must be applied to single propositions (or two-element partitions) only.

The question that now needs to be addressed is whether linear averaging in response to opinions of others on a *single* proposition is compatible with Bayesian conditioning. To answer this question we will restrict ourselves to the following kind of situation. You and n experts are Bayesian reasoners in the sense of having probabilistic degrees of belief for all propositions of interest, which are revised by conditionalisation on any evidence that is acquired. (In some of the examples to be considered, You are one of the experts.) Let P: $\mathcal{S} \rightarrow [0,1]$ be your prior probability for these propositions and $P_1, ..., P_n$ be the priors of the experts, with domains that overlap to some relevant extent with $\mathcal{S}$. (To avoid problems with zero probabilities, let's assume that these priors assign non-zero probability to any proposition that is conditionalised upon.) Each expert I acquires some evidence, the proposition E_i, and having conditioned on it, finds herself with posterior probabilistic beliefs R_i. What proposition she has learnt is unknown to You and the other experts, but each expert reports her new partial beliefs on some particular proposition X. Your revised probabilities, after having heard the experts' reports is given by a posterior probability Q. The focus of our interest is the form that Q must take if You revise your opinions rationally in response to what the experts say.

Let us consider first what is required of You by Bayesianism. If you are to revise by conditionalisation then the domain $\mathcal{S}$ of your beliefs must contain propositions identifying the possible reports by the experts. Risking some ambiguity in the interests of economy, let $\mathcal{R}_i$ be a random variable ranging over posterior opinion states of expert I and let R_i denote the proposition that expert I's posterior probability is R_i (i.e. R_i serves as shorthand for $\mathcal{R}_i = R_i$, as well as denoting I's posterior probabilities). Similarly, let $R_i(X)$ be a random variable ranging over possible reports on X by expert I and let $R_i(X) = x_i$ be a proposition expressing the fact that expert I's reported posterior probability for X is x_i. Then in response to this report a supra-Bayesian forms a new belief in the proposition X of:

$$Q(X) = P(X|R_i(X) = x_i)$$

Now these reports can serve as opinion functions on the domain $\{X\}$. And a corresponding opinion formation rule F on this domain can be said to be *Bayes-compatible* just in case for any profile of reports $R_1(X) = x_i$, ..., and $R_n(X) = x_n$ and any common prior P:

$$P(X|R_1(X) = x_1, ..., R_n(X) = x_n) = F(R_1, ..., R_n)(X) \tag{1}$$

In particular linear averaging on this domain will be Bayes-compatible iff for some weights α_i such that $\alpha_i \geq 0$ and $\sum_i \alpha_i = 1$:

$$P(X|R_1(X) = x_1, ..., R_n(X) = x_n) = \sum_{i=1}^{n} \alpha_i.R_i(X) \tag{2}$$

So our question boils down to whether it is always possible to find proposition-dependent weights on the experts' opinions on X such that equation 2 is satisfied. A positive answer would not in itself show that linear averaging on one proposition at a time is a sensible way of revising one's beliefs, but it would at least show that it was not obviously unsuitable. In contrast if, for a range of typical cases, no such weights could be found, then we can conclude that linear averaging is not a candidate for a general method of belief revision in response to testimonial evidence.

3 Full Deference

As many authors have observed, we might defer to someone's opinion because we think that they hold information that we do not or because we believe them to have skills which make them better at judging the significance of the information that we both hold (see for instance, Joyce [13] and Elgar [8]). In this section we look at what such deference amounts to in these two cases, focusing for the moment on cases of the first kind.

3.1 Deferring to Information

Let us begin with a very simple situation, involving just one expert and where You and the expert share a common prior P over some domain. Suppose the expert acquires evidence E and as a result adopts posterior R. Because she is strictly better informed than You and because you share a common prior, You should defer to the expert's opinion in the sense of satisfying:

Global Deference: $P(\cdot|\mathcal{R} = R) = R(\cdot)$

Global Deference implies that your prior probability for any proposition equals the expected opinion of the expert on it, i.e. that $\mathbb{E}(R(\cdot)) = P(\cdot)$.[2] That is, before You have received the expert's report, You expect her to report what You currently believe, because what You believe is based on all the evidence that You have at that time. Or to put it the other way round, if You believed the expert to hold evidence regarding the truth of X that makes it more probable than You judge it to be, then You should immediately adjust your belief in X. Evidence of evidence for X is evidence for X.

A Bayesian agent who defers to the expert's superior information in the sense of respecting Global Deference is disposed to adopt the expert's posterior opinions R as her own when she learns what these are. In the envisaged circumstances however it is not possible to infer all the expert's opinions simply from her report on a single proposition. Nonetheless, on hearing the expert's report on proposition X, You should conform with Global Deference to the extent of adopting the expert's opinion on X by satisfying:

[2] See Joyce [13, p. 191] for a generalisation of this claim.

Local Deference: $Q(X) = P(X|R(X) = x) = P(X|E) = x$

In other words, your revised probability for X should be the probability of proposition X given the expert's report (in accordance with Bayesianism), which is just the probability of X given the evidence obtained by the expert (in virtue of the common prior), which is just the probability for X reported by her (in virtue of her truthfulness).

Now conformity to Local Deference does not determine your posterior belief state; just a fragment of it. To propagate the implications of deferring to the expert's reported opinion on X to the rest of one's beliefs, one might reasonably follow Steele's [18] recommendation to Jeffrey conditionalise on the partition $\{X, \neg X\}$ taking as inputs one's newly revised degrees of belief for its elements. This yields, for all propositions $Y \in \mathcal{S}$:

$$\begin{aligned} Q(Y) &= P(Y|X).Q(X) + P(Y|\neg X).Q(\neg X) \\ &= P(Y|X).x + P(Y|\neg X).(1 - x) \end{aligned}$$

Revising your beliefs in this fashion is demonstrably the rational thing to do just in case You take the evidential significance for Y of the expert's reports on X to be screened out by the truth of X or $\neg X$, i.e. just in case $P(Y|R(X) = x, X) = P(Y|X)$ and $P(Y|R(X) = x, \neg X) = P(Y|\neg X)$. For then:

$$\begin{aligned} Q(Y) &= P(Y|R(X) = x) \\ &= P(Y|X, R(X) = x).P(X|R(X) = x) + P(Y|\neg X, R(X) = x).P(\neg X|R(X) = x) \\ &= P(Y|X).x + P(Y|\neg X).(1 - x) \end{aligned}$$

A couple of cautionary points. Firstly, it must emphasised that revising one's beliefs in this fashion, namely by application of Local Deference and Jeffrey conditionalisation, does not guarantee that one's posterior beliefs will be the same as the expert's, i.e. that for all Y in the belief domain, $Q(Y) = P(Y|E)$. (In the absence of further information, however, one will expect them to be so.)

Secondly, one cannot mechanically apply this type of revision to a sequence of reports by the expert. Suppose, for instance, having deferred to the expert's opinion on X and then Jeffrey conditioned on the partition $\{X, \neg X\}$, the expert reports her opinion on Y. It is clear that You should defer to this opinion as well, since the expert is still strictly better informed than You. But if You attempt once again to revise your other beliefs by Jeffrey conditioning, this time on the partition $\{Y, \neg Y\}$, using your newly acquired opinions on its members as inputs, You will be led to revise your opinion on X (except, of course, when X and Y are probabilistically independent). And this would conflict with your commitment to adopting the expert's opinion on X as your own. What You should do is revise your other beliefs subject to the *dual* constraint on the partition $\{XY, X\neg Y, \neg XY, \neg X\neg Y\}$ implied by your deference to the expert's reported opinions on both X and Y. However, this constraint does not determine a unique redistribution of probability across the relevant partition. So we have not settled the question of how, in general, one should respond to multiple reports by a single expert.

3.2 Experts with Different Priors

Let us now drop the assumption of common priors and look at cases in which the expert has special skills rather than additional information. Suppose for example that an investor has doubts about the financial viability of a company in which she holds shares. The investor gets access to the company's accounts but, lacking expertise in accounting, has the books examined by an accountant. The accountant looks over the books and declares the company to be in a sound financial situation. Although the investor may completely trust the accountant's judgement of the evidence provided by the company's accounts, she may nonetheless not be willing to simply adopt the accountant's posterior judgement, because she thinks that the accountant's prior was based on less evidence than her own.

In Bayesian statistics the support that some piece of evidence gives to one hypothesis relative to another is often measured, in a prior-free way, by its *Bayes factor*.[3] We arrive at this factor in the following way. Note that for any probability P and evidence proposition E:

$$\frac{P(X|E)}{P(Y|E)} = \frac{P(E|X)}{P(E|Y)}.\frac{P(X)}{P(Y)}$$

So let the associated Bayes factor, $\mathcal{B}(X,Y)$, for proposition X relative to proposition Y, be defined by:

$$\mathcal{B}(X,Y) := \frac{P(E|X)}{P(E|Y)}$$

Then it follows that:

$$\frac{P(X|E)}{P(Y|E)} = \mathcal{B}(X,Y).\frac{P(X)}{P(Y)} \tag{3}$$

Now in situations in which one wishes to defer to an expert's judgement on the significance of some evidence E, but not to their prior opinions, one can use the decomposition given by equation 3 to determine one's posterior probabilities from the expert's Bayes factors and one's own prior opinion. To illustrate, suppose, as before, that the expert reports both their prior and posterior probabilities for the proposition X. From this their Bayes factor

[3] See Jeffreys [12] for the classic statement of the case for this measure.

for X relative to $\neg X$ can be determined.[4] Then You should defer to the expert's judgement on the significance of E for X by setting your posterior probability for X to satisfy the following principle.

Bayes Deference: $Q(X) = P(X|\beta(X, \neg X) = b) = \frac{b.P(X)}{b.P(X)+P(\neg X)}$

Revision of belief by application of Bayes Deference to a single proposition does not determine the impact of the expert's testimony on your other beliefs. But, with the same caveats as before, we may again combine it with Jeffrey conditionalisation, taking as inputs the values $Q(X)$ and $Q(\neg X)$, in order to propagate the change to all other propositions.

4 Partial Deference

4.1 Single Expert

For Bayesian agents who revise their beliefs by conditionalisation, adherence to Local Deference amounts to a disposition to adopt whatever opinions on X that the expert reports, a disposition that is rationalised in the circumstances under consideration by the fact that she holds more information. Linear averaging is compatible with Bayesian conditionalisation in these circumstances if and only if full weight is given to the expert's opinion. For your revised probabilities, Q, after hearing the expert's report on X, are required to satisfy both:

$$Q(X) = R(X)$$
$$Q(X) = \alpha.R(X) + (1 - \alpha).P(X)$$

and this implies that $\alpha = 1$. So in situations in which full deference is appropriate, supra-Bayesian belief revision is trivially consistent with linear averaging.

Situations in which full deference is appropriate are obviously quite special however. Other situations call for only *partial* deference. For instance when You and the expert both acquire information, then the expert's information is no longer strictly greater than yours and deferring to their opinion would be tantamount to discarding the additional information you

[4] For:

$$\begin{aligned}\mathcal{B}(X, \neg X) &= \frac{P(E|X)}{P(E|\neg X)} \\ &= \frac{P(X|E)}{P(\neg X|E)}.\frac{P(\neg X)}{P(X)} \\ &= \frac{R(X)}{1 - R(X)}.\frac{1 - P(X)}{P(X)}\end{aligned}$$

hold. Suppose that You and the expert have respectively acquired new information E and E_1 and adopted posterior beliefs $P(\cdot|E)$ and $R = P(\cdot|E_1)$. Now You should defer, not to the expert's posterior opinions, but to her *conditional* opinions, given the information E that You have acquired. This means that we apply Global Deference, not to P and R, but to $P(\cdot|E)$ and $R(\cdot|E)$, the latter pair being based on strictly more information than the former pair. This yields:

Conditional Deference:

$$P(\cdot|\mathcal{R} = R, E) = R(\cdot|E) \tag{4}$$

Equation 4 gives a Bayesian expression for partial deference to an expert appropriate to the circumstances of shared priors but different information. It does not suffice however to determine how You should revise your beliefs in response to the expert's report on X, since it is her conditional opinion on X given that E that You wish to adopt, not her expressed unconditional opinion on X. So there is now an additional motive to ask whether the kind of partial deference to the expert's opinion represented by linear averaging can further constrain your opinion in a Bayes-compatible way. That is, we want to know whether You can assign a weight α to the expert's opinion, expressing your degree of deference, such that:

$$P(X|R_1(X), E) = \alpha.R_1(X) + (1-\alpha).P(X|E)$$

Instead of tackling this question directly, let us just note that the situation is, formally speaking, very similar one in which You can consult two experts, each of which has acquired information over and above that held by You. In such a case it would also not be appropriate to defer completely to any one of the experts, as this would be tantamount to discarding the information acquired by the others. Indeed, the first case can be treated a special instance of two-expert ones, by letting one of the two experts be You. So we can study the two cases together.

4.2 Multiple Experts

Suppose that two experts, again sharing prior P with You, have respectively acquired new information E_1 and E_2 and adopted posterior beliefs R_1 and R_2. In this case, in virtue of the shared priors, Global Deference tells us that:

$$P(\cdot|\mathcal{R}_1 = R_1) = R_1(\cdot) = P(\cdot|E_1) \tag{5}$$

$$P(\cdot|\mathcal{R}_2 = R_2) = R_2(\cdot) = P(\cdot|E_2) \tag{6}$$

And from these two equations it follows that:

$$P(\cdot|\mathcal{R}_1 = R_1, \mathcal{R}_2 = R_2) = P(\cdot|E_1, E_2) \tag{7}$$

But since You don't know what the propositions E_1 and E_2 are that the experts have learnt, this is again not enough to determine how You should revise your beliefs in X. Nonetheless we can now ask whether linear averaging is consistent with equations 5, 6 and 7, i.e., whether for each proposition X there exists a weight α such that:

$$P(X|R_1(X), R_2(X)) = \alpha.R_1(X) + (1-\alpha).R_2(X) \tag{8}$$

where, as indicated before, α can depend on X but not on what the experts report about X.

The answer is that no such weight exists. The reason for this is quite simple: Local Deference together with equation 8 requires that the weight You put on each expert's report *does* depend on what they report. To see this consider the following example. Suppose that $R_1(X) = x_1$ and $R_2(X) = 0$. Then by Local Deference $P(X|R_1(X) = x_1) = x_1$ and $P(X|R_2(X) = 0) = 0$. It follows that:

$$P(X|R_1(X) = x_1, R_2(X) = 0) = 0$$

So if You conditionalise on the experts' reports on X then your posterior beliefs are such that $Q(X) = 0$. But by linear averaging:

$$\begin{aligned} Q(X) &= \alpha.R_1(X) + (1-\alpha).R_2(X) \\ &= \alpha.x_1 + (1-\alpha).0 \end{aligned}$$

and this implies that $\alpha = 0$. But by the same argument if $R_1(X) = 0$ and $R_2(X) = x_2$ then for Bayes-compatibility of linear averaging, it must be the case that $\alpha = 1$. So the weights on the experts' reports on X are *not* independent of what they report.

We can put the point slightly differently. If the weights on the experts' reports on X are assumed to be independent of what they report, then we can infer that they always make the same report. For in our example we assumed that the experts made different reports in order to infer weight variability. But if weights can't vary then the experts cannot have made different reports. It is true that this example made essential use of the fact that experts can report zero probabilities for propositions, but restricting linear averaging to cases in which they report non-zero probabilities makes no difference to the conclusion. For Global Deference implies that You expect, conditional on expert 1 reporting a probability of x for X, that your posterior degree of belief in X, after hearing *both* experts' reports, to be just x, i.e. that:

$$\mathbb{E}(Q(X)|R_1(X) = x) = x \tag{9}$$

(This is proved as Theorem 1 in the appendix). But equation 9, together with assumption that your posterior beliefs are a linear average of the experts', implies that your conditional expectation for expert 2's report on X, given that expert 1 reports a probability of x for X, equals x as well (Theorem 2 in the appendix). And this can be the case only if the random variables corresponding to each of the expert's opinion reports are the same, i.e. if $\mathcal{R}_1 = \mathcal{R}_2$ (Corollary 1). So, on pain of contradiction, your posterior opinion on X can be both a linear average of the expert's reported opinions and defer appropriately to these reports only if they are same.

4.3 Concluding Remarks

We have shown that revising your beliefs on some proposition X by taking the linear average of reported expert opinion is consistent with Bayesian conditioning, as guided by the principle of Global Deference, only in the trivial case when the experts always make the same report (with probability one). This I take to rule linear averaging out as a rational response to disagreement in expert opinion. It also therefore rules out linear averaging as a response to disagreement with your epistemic peers, construed as others who share your priors but hold different information. For this case is formally equivalent to that of disagreement amongst experts with different information.

A couple of cautionary notes about the scope of these conclusions. Firstly, nothing has been said in this paper about other forms of averaging or indeed forms of opinion pooling that do not involve averaging. Guidance on this question can be found in Dawid, DeGroot and Mortera [6], where a much more extensive set of formal results on the Bayes-compatibility of opinion pooling is proved. The upshot of these results is far from settled, but it would seem that while log-linear and geometric averaging in their usual forms are subject to the same difficulties as linear averaging, more general versions of all of these rules are not so easily dismissed. The philosophical status of these more general rules deserve further exploration.

Secondly, it should also be emphasised that this study leaves open the question of whether linear averaging is the appropriate response to situations in which you find yourself in disagreement with peers who hold the same information as you and are as a good at judging its significance. In the philosophical literature, the view that one should respond to such disagreements by taking an equal-weighted average of your opinions has been hotly debated. But nothing presented here militates either for or against this view.

5 Appendix: Proofs

Lemma 1 *Suppose that* $\mathbb{E}(X|Y = y) = y$ *and* $\mathbb{E}(Y|X = x) = x$*. Then* $X = Y$*.*

Proof Note that by the law of iterated expectations, $\mathbb{E}(Y) = \mathbb{E}(\mathbb{E}(X|Y)) = \mathbb{E}(X)$. Now since $(Y|X = x) = x$, it follows by another application of the law of iterated expectations that:

$$\begin{aligned}\mathbb{E}(X\cdot Y) &= \mathbb{E}(\mathbb{E}(XY|X))\\ &= \mathbb{E}(X\cdot\mathbb{E}(Y|X))\\ &= \mathbb{E}(X^2)\end{aligned}$$

It follows that:

$$\begin{aligned}cov(X,Y) &= \mathbb{E}[(X-\mathbb{E}(X))(Y-\mathbb{E}(Y)]\\ &= \mathbb{E}(X\cdot Y)-\mathbb{E}(X)\cdot\mathbb{E}(Y)\\ &= \mathbb{E}(X^2)-[\mathbb{E}(X)]^2\end{aligned}$$

since $\mathbb{E}(X\cdot Y)=\mathbb{E}(X^2)$ and $\mathbb{E}(Y)=\mathbb{E}(X)$. So $cov(X,Y)=var(X)=var(Y)$. Hence $var(X-Y)=var(X)+var(Y)-2cov(X,Y)=0$. So $X=Y$ with probability 1.

Theorem 1 *Assume Global Deference with respect to each expert i, i.e. that for $i=1,2$, $P(X|R_i(X))=R_i(X)$. Then $\mathbb{E}(Q(X)|R_i(X)=\bar{x}_i)=\bar{x}_i$*

Proof By Deference, $P(X|R_1(X)=\bar{x}_1)=\bar{x}_1$. And:

$$\begin{aligned}P(X|R_1(X)=\bar{x}_1) &= \textstyle\sum_i P(X,R_2(X)=x_i|R_1(X)=\bar{x}_1))\\ &= \textstyle\sum_i P(X|R_2(X)=x_i,R_1(X)=\bar{x}_1))\cdot P(R_2(X)=x_i|R_1(X)=\bar{x}_1)\\ &= \mathbb{E}(Q(X)|R_1(X)=\bar{x}_1)\end{aligned}$$

Hence $\mathbb{E}(Q(X)|R_1(X)=\bar{x}_1)=\bar{x}_1$. Similarly, $\mathbb{E}(Q(X)|R_2(X)=\bar{x}_2)=\bar{x}_2$.

Theorem 2 *Assume Global Deference with respect to expert 1 and 2 and that for $0<\alpha<1$, $Q(X)=\alpha\cdot R_1(X)+(1-\alpha)\cdot R_2(X)$. Then:*

$$\begin{aligned}\mathbb{E}(R_1(X)|R_2(X)=\bar{x}_2)=\bar{x}_2\\ \mathbb{E}(R_2(X)|R_1(X)=\bar{x}_1)=\bar{x}_1\end{aligned}$$

Proof By Theorem 1, $\mathbb{E}(\alpha\cdot R_1(X)+(1-\alpha)\cdot R_2(X)|R_1(X)=\bar{x}_1)=\bar{x}_1$. Hence by the Linearity of expectations:

$$\begin{aligned}\alpha\cdot\mathbb{E}(R_1(X)|R_1(X)=\bar{x}_1))+(1-\alpha)\cdot\mathbb{E}(R_2(X)|R_1(X)=\bar{x}_1))\\ =\alpha\cdot\bar{x}_1+(1-\alpha)\cdot\mathbb{E}(R_2(X)|R_1(X)=\bar{x}_1)\\ =\bar{x}_1\end{aligned}$$

But this can be the case iff $\mathbb{E}(R_2(X)|R_1(X)=\bar{x}_1)=\bar{x}_1$. Similarly $\mathbb{E}(R_1(X)|R_2(X)=\bar{x}_2)=\bar{x}_2$.

Corollary 1 $\mathcal{R}_1=\mathcal{R}_2$

Proof Follows from Theorem 2 and Lemma 1.

Acknowledgements This paper has benefited from presentation at a workshop on 'Aggregation, Deliberation and Consensus' at the DEC, Ecole Normale Superieure, organised by Mikaël Cozic and one at the PSA organised by Mikaël Cozic and Jan-Willem Romeijn. Franz Dietrich, Wichers Bergsma and Angelos Dassios helped with Lemma 1. I am grateful to the editors of the special edition and an anonymous referee for their comments on an earlier draft.

References

1. Bradley, R. (2006) “Taking Advantage of Difference of Opinion”, *Episteme* 3: 141-155
2. Bradley, R. (2007) “Reaching a Consensus", *Social Choice and Welfare* 29: 609-632
3. Bonnay, D. & Cozic, M., (2015b), “Weighted averaging as a revision rule”, manuscript
4. Bonnay, D. & Cozic, M., (2015c), “Jeffrey conditioning and Weighted averaging”, manuscript
5. Bradley, R. (2007) “The Kinematics of Belief and Desire”, *Synthese* 156: 513-535
6. Dawid, A.P., M.H. DeGroot and J. Mortera (1995) "Coherent Combination of Experts’ Opinions", *Test* 4: 263-313
7. Dietrich, F. and C. List (forthcoming) "Probabilistic opinion pooling", Oxford Handbook of Probability and Philosophy
8. Elgar, A. (2007) “Reflection and disagreement”, *Noûs* 41 (3):478–502
9. French, S. (1981) "Consensus of opinion". Eur. J. Oper. Res. 7: 332-340.
10. Genest, C. and M. J. Schervish "Modeling Expert Judgments for Bayesian Updating", *Annals of Statistics* 13 (3): 1198-1212.
11. Genest, C. and J. V. Zidek (1986) “Combining Probability Distributions: A Critique and Annotated Bibliography”, *Statistical Science* 1: 113-135
12. H. Jeffreys (1961). *The Theory of Probability* (3 ed.). Oxford: Oxford University Press
13. Joyce, J. (2007) "Epistemic Deference: The Case of Chance," *Proceedings of the Aristotelian Society* 107: 1-20.
14. Lehrer, K. (1976) “When Rational Disagreement is Impossible”, *Noûs* 10: 327-332
15. Lehrer, K and C. Wagner (1981) *Rational Consensus in Science and Society*, Dordrecht: Reidel
16. Morris, P. A. (1974) "Decision analysis expert use", Management Science 20: 1233-1241.
17. Romeijn, J-W. (2015) "Pooling, Voting and Bayesian Updating", manuscript
18. Steele, K. (2012) "Testimony as Evidence", *Journal of Philosophical Logic* 41: 983-999
19. Wagner, C. (1978) “Consensus through Respect: A Model of Rational Group Decision-making”, *Philosophical Studies* 34: 335-349
20. Wagner, C. (1982) “Allocation, Lehrer Models, and the Consensus of Probabilities”, *Theory and Decision* 14: 207-220
21. Wagner, C. (1985) “On the Formal Properties of Weighted Averaging as a Method of Aggregation”, *Synthese* 62: 97-108