

On the accuracy of group credences*

Richard Pettigrew
Richard.Pettigrew@bris.ac.uk

October 8, 2016

We often ask for the opinion of a group of individuals. How strongly does the scientific community believe that the rate at which sea levels are rising increased over the last 200 years? How likely does the Monetary Policy Committee of the Bank of England think it is that there will be a recession if the country leaves the European Union? How confident is the scholarly community that William Shakespeare wrote *Hamlet*? Suppose you ask me one of these questions, and I respond by listing, for each member of the group in question, the opinion that they hold on that topic. I list each scientist in the scientific community, for instance, and I give their credence that sea level rise accelerated in the past two centuries. By doing this, I may well give you enough information so that you can calculate the answer to the question that you asked; but what I give you does not amount to that answer. What you were asking for was not a set of credences, one for each member of the group; you were asking for a single credence assigned collectively by the group as a whole. What is this group credence? And how does it relate to the individual credences assigned by the members of the group in question?

In this paper, I'd like to explore a novel argument for a familiar partial answer to the latter question. In particular, given a group of individuals, I'd like to say that any account of how we aggregate the credences of those individuals to give the credences of the group must have a particular property — the group credences should be a weighted average of the individual credences. Now, a weighted average is sometimes called a *linear pool*, and this constraint on the aggregation of credences is usually called *linear pooling*. I will not have much to say about how we should set the weightings when we first take our linear pool of the individual credences. But I will have something to say about how those weightings should evolve as new evidence arrives. I will also have something to say about the two standard objections to linear pooling as a constraint on the aggregation of credences.

1 Group opinions and group knowledge

Before I present my argument, I'd like to say a little more about the first question from above. This was the more ontological of the two. It asked: What *is* a group credence? Indeed, it

***Acknowledgments:** I am very grateful to a number of people for helpful feedback on earlier versions of this paper: Alexander Bird, Liam Kofi Bright, Catarina Dutilh Novaes, Branden Fitelson, Brian Hedden, Remco Heesen, Matt Kopec, Jan-Willem Romeijn, Jeff Sanford Russell, Julia Staffel, Brian Weatherson, Greg Wheeler, Kevin Zollman, as well as two anonymous referees for this journal. However, the paper reports my credences, not those of that group — I take responsibility for any flaws. This work was supported by the European Research Council Seventh Framework Program (FP7/2007-2013), ERC grant 308961-EUT.

turns out that there are at least two different notions that might go by that name. I will be concerned with the notion of a group credence function at a time as a summary of the credal opinions of the individuals in the group at that time. But we might also think of the notion of a group credence function at a time as a summary of the potential knowledge that is distributed throughout the credal opinions of the individuals in the group at that time.

Let's see these different notions in action. Suppose two historians, Jonathan and Josie, are researching the same question, but in two different archives. Both know that there may be a pair of documents, one in each archive, whose joint existence would establish a controversial theory beyond doubt. Jonathan finds the relevant document in his archive, but doesn't know whether Josie has found hers; and Josie finds the relevant document in her archive, but doesn't know whether Jonathan has found his. Indeed, each assigns a very low credence to the other finding their document; as a result, both have a very low credence in the controversial theory. According to the first notion of a group credence function as a summary of the credal opinions of the group's members, the group credence in the controversial theory should remain low. After all, both members of the group assign it a low credence. However, according to the second notion of a group credence function as a summary of the knowledge that is distributed throughout the credal opinions of the individuals in the group, the group credence in the controversial theory should be high. After all, between them, they know that both documents exist, and they both agree that this would give extremely strong evidence for the controversial hypothesis.

These two notions of a group credence function are not independent. It seems natural to think that we arrive at the second sort of group credence function — the sort that summarises the knowledge distributed throughout the group — as follows: we take the credences of each of the members of the group; we let them become common knowledge within the group; we allow the members of the group to update their individual credences on this common knowledge; and then, having done this, we take the first notion of a group credence function — that sort that summarizes the credal opinions of the individuals in the group — with these updated individual credence functions as its input. That is, group credences as summaries of distributed potential knowledge are simply summaries of the opinions of the individuals in the group *once those individuals are brought up to speed with the opinions of the other individuals*.¹ Thus, once Jonathan learns that Josie has very high credence that the relevant document from her archive exists, and once Josie learns that Jonathan has very high credence that the relevant document from his archive exists, they will both update to a high credence that both documents exist, and from there to a high credence that the controversial theory is true. Taking a summary of those updated credences gives us a high group credence that the controversial theory is true — and of course that is the correct summary of the potential knowledge distributed between Jonathan and Josie.

Throughout, I'll be interested in the first sort of group credence function. That is, I'll be concerned with defending a norm for group credences understood as summaries of the

¹Of course, much thought has been given to how you should respond when you learn of someone else's credences; this is the focus of the peer disagreement literature. For accuracy-based analyses of this problem, see (Moss, 2011; Staffel, 2015; Levinstein, 2015; Heesen & van der Kolk, 2016). Also, if the members of the group share the same prior probabilities, Robert Aumann's famous Agreement Theorem shows that there is just one rational way for members of the group to respond—they must all end up with the same credence function once they have updated on the common knowledge of one another's posterior probabilities (Aumann, 1976). Thus, in such a case, there will be no disagreement to resolve. But of course there are many cases in which things are not so simple, because the members of the group do not share the same priors.

opinions of the individuals in the group.

2 Linear pooling

Having pinned down the notion of group credence function that will concern us here, let me introduce an example to illustrate the account of group credence functions that I will defend here, namely, linear pooling. Adila, Benicio, and Cleo are climate scientists. They all have opinions about many different propositions; but all three have opinions about the following proposition and its negation: *The sea level rise between 2016 and 2030 will exceed 20cm*. We'll call this proposition H ; and we'll write \bar{H} for its negation. The following table gives the credences that Adila (cr_A), Benicio (cr_B), and Cleo (cr_C) assign to the two propositions:

	cr_A	cr_B	cr_C
H	0.2	0.4	0.5
\bar{H}	0.8	0.6	0.5

But what credences does the group, Adila-Benicio-Cleo, assign to these two propositions? There are two standard proposals in the literature: *linear pooling* and *geometrical pooling*. On both, each member of the group is assigned a real number as a weighting—thus, we might let α_A be Adila's weighting, α_B for Benicio, and α_C for Cleo. We assume that these weightings are non-negative real numbers; and we assume that they sum to 1.

According to linear pooling, the group credence in a particular proposition is given by the relevant weighted arithmetic average of the individual credences in that proposition. For instance,

$$Cr_{LP}(H) = \alpha_A cr_A(H) + \alpha_B cr_B(H) + \alpha_C cr_C(H) = 0.2\alpha_A + 0.4\alpha_B + 0.5\alpha_C$$

So, for example, if each weighting is the same, so that $\alpha_A = \alpha_B = \alpha_C = \frac{1}{3}$, then

$$Cr_{LP}(H) = (0.2 \times \frac{1}{3}) + (0.4 \times \frac{1}{3}) + (0.5 \times \frac{1}{3}) \approx 0.37$$

and

$$Cr_{LP}(\bar{H}) = (0.8 \times \frac{1}{3}) + (0.6 \times \frac{1}{3}) + (0.5 \times \frac{1}{3}) \approx 0.63$$

More generally:

Linear Pooling Suppose:

- G is a group of individuals;
- the credence functions of its members are cr_1, \dots, cr_n ;
- cr_i is defined on the algebra of propositions \mathcal{F}_i ;
- the credence function of the group is Cr_G ;
- Cr_G is defined on the algebra of propositions $\mathcal{F} = \bigcap_{i=1}^n \mathcal{F}_i$.

Then there should be real numbers $\alpha_1, \dots, \alpha_n \geq 0$ with $\sum_{i=1}^n \alpha_i = 1$ such that

$$Cr_G(-) = \sum_{i=1}^n \alpha_i cr_i(-)$$

If this holds, we say that Cr_G is a *weighted (arithmetic) average* or *mixture* or *linear pool* of cr_1, \dots, cr_n .

Notice that, if the individuals' credence functions are probability functions, then so is any linear pool of them. We will assume throughout that the individuals' credence functions are indeed probability functions. And we will assume that each of the set of propositions \mathcal{F}_i on which cr_i is defined is finite — thus, $\mathcal{F} = \bigcap_{i=1}^n \mathcal{F}_i$ is also finite.

What's so bad about a group credence function that is not a mixture of the individual credence functions held by the members of the group? That is, what is so bad about violating Linear Pooling? And what's so good about Cr_{LP} ? Roughly, our answer will be based on the following pair of claims:

- (I) If Cr_G is not a mixture of cr_1, \dots, cr_n , then there is an alternative group credence function Cr_G^* such that, by the lights of each member of the group, the expected epistemic value of Cr_G^* is *strictly greater* than the expected epistemic value of Cr_G
- (II) If Cr_G is a mixture of cr_1, \dots, cr_n , then there is not even an alternative group credence function Cr_G^* whose expected epistemic value is *at least* that of Cr_G by the lights of each member of the group.

Consider, for instance, the candidate group credence function Cr_{ABC} that assigns 0.1 to H and 0.9 to \bar{H} . It is not a linear pool of cr_A , cr_B , and cr_C . What is wrong with it? Well, according to this argument, the following is true: there is some alternative group credence function Cr^* that Adila, Benicio, and Cleo all expect to do better than Cr . And Cr^* is itself a linear pool of cr_A , cr_B , and cr_C . And there is no alternative group credence function Cr' that Adila, Benicio, and Cleo all expect to do better than Cr^* . In the next few sections, we make this more precise. We start, in section 3, by making precise how we measure epistemic value. Then, in section 4, we make precise why the claim just stated would, if true, establish Linear Pooling as a rational requirement; and we observe that a precise version of that claim is true. Finally, in sections 5 and 6, we address two common objections to Linear Pooling — it does not preserve probabilistic independences; and it does not commute with Bayesian conditionalization.

3 Epistemic value and accuracy

The account of epistemic value that I favour is monistic. That is, I think there is just one fundamental source of value for credences. Following James M. Joyce (1998) and Alvin Goldman (2002), I take the epistemic value of a credence to be its accuracy. A credence in a true proposition is more accurate the higher it is; a credence in a false proposition is more accurate the lower it is. Put another way: the ultimate goal of having credences is to have maximal credence (i.e. credence 1) in truths and minimal credence (i.e. credence 0) in falsehoods; the epistemic value of a credence is given by its proximity to that ultimate goal. This is a credal version of veritism.

How should we measure the accuracy of a credence? In fact, in keeping with other writings in this area, we will talk about measuring the *inaccuracy* of a credence rather than its accuracy; but I take the accuracy of a credence to be simply the negative of its inaccuracy, so there is nothing substantial about this choice. Now, the inaccuracy of a credence in a proposition should depend only on the credence itself, the proposition to which it is assigned, and the truth value of the proposition at the world where its inaccuracy is being assessed. Thus, we can measure the inaccuracy of a credence using what has come to be called a *scoring rule*.

A scoring rule s is a function that takes a proposition X , the truth value i of X (represented numerically, so that 0 represents *falsity* and 1 represents *truth*), and the credence x in X that we will assess; and it returns a number $s_X(i, x)$. So $s_X(1, x)$ is the inaccuracy of credence x in proposition X when X is true; and $s_X(0, x)$ is the inaccuracy of credence x in proposition X when X is false.

What features should we require of a scoring rule s if it is to be a legitimate measure of inaccuracy? We will consider two.

First, we will require that the inaccuracy of a credence varies continuously with the credence. That is:

Continuity $s_X(1, x)$ and $s_X(0, x)$ are continuous functions of x .

I take this to be a natural assumption.

Second, we will require that the scoring rule s should be *strictly proper*. Take a particular credence p in a proposition X — so p is a real number at least 0 and at most 1. We can use p to calculate the expected inaccuracy of any credence x , including p itself. The expected inaccuracy of x from the point of view of p is simply this:

$$p s_X(1, x) + (1 - p) s_X(0, x)$$

That is, it takes the inaccuracy of x when X is true — namely, $s_X(1, x)$ — and weights it by p , since p is a credence in X ; and it takes the inaccuracy of x when X is false — namely, $s_X(0, x)$ — and weights it by $(1 - p)$, which is of course the corresponding probabilistic credence in the negation of X . A scoring rule is strictly proper if, for any such credence p in a proposition X , it expects itself to have lowest inaccuracy out of all the possible credences in X . That gives us the following principle:

Strict Propriety For any proposition X and any $0 \leq p \leq 1$,

$$p s_X(1, x) + (1 - p) s_X(0, x)$$

is minimized uniquely, as a function of x , at $p = x$.

This demand has been justified in various ways: see (Gibbard, 2008), (Joyce, 2009), or (Pettigrew, 2016, Chapters 3 and 4). Let me briefly sketch Joyce's argument for Strict Propriety. Suppose you are a veritist about credences: that is, you hold that the sole fundamental source of epistemic value for credences is their accuracy; and you hold that facts about the rationality of a credal state are determined by facts about its epistemic value. Then reason as follows: Let X be a proposition and p a credence. Then, intuitively, there is evidence you might receive to which the unique rational response is to set your credence in X to p . For instance, you might learn with certainty that the objective chance of X is p . Or your most revered epistemic guru might whisper in your ear that she has credence p in X . Suppose you do receive this evidence; and suppose that you do set your credence in X to p , in accordance with your evidence. Now, suppose that, contrary to Strict Propriety, there is some other credence $x \neq p$ such that p expects x to be at most as inaccurate as p expects p to be: that is, $p s_X(1, x) + (1 - p) s_X(0, x) \leq p s_X(1, p) + (1 - p) s_X(0, p)$. Then, from the point of view of veritism, there would be nothing epistemically reprehensible were you to switch from your current credence p in X to a credence of x in X without obtaining any new evidence in the meantime. After all, we typically think that it is rationally permissible to adopt an option

that you currently expect to be at most as bad as your current situation. And, by hypothesis, p expects credence x in X to be at most as inaccurate — and thus, by veritism, at most as bad, epistemically speaking — as it expects credence p in X to be. But of course making such a switch would be epistemically reprehensible. It is not rationally permissible for you to shift your credence in X from p to x without obtaining any new evidence in the meantime; for, by hypothesis, having credence p in X is the *unique* rational response to your evidence.² Therefore, there can be no such credence x in X . Thus, in general, for any credence p in any proposition X , there can be no alternative credence $x \neq p$ in X that p expects to be at most as inaccurate as it expects itself to be. And that is just what it means to say that the measure of inaccuracy must be strictly proper.

What do these strictly proper and continuous scoring rules look like? Here is an example. It is called the *quadratic scoring rule*:

- $q_X(1, x) := (1 - x)^2$
- $q_X(0, x) := x^2$

In other words: $q_X(i, x) = |i - x|^2$, where $i = 0$ or 1 . So the quadratic scoring rule takes the difference between the credence you have and the credence you would ideally have — credence 1 if the proposition is true and credence 0 if it's false — and squares that difference to give your inaccuracy. A little calculus shows that q is strictly proper and continuous.

Here is an example of a scoring rule that is continuous, but not strictly proper. It is called the *absolute value measure*:

- $a_X(1, x) := 1 - x$
- $a_X(0, x) := x$

In other words: $a_X(i, x) = |i - x|$. So the absolute value measure takes the difference between the credence you have and the credence you would ideally have and takes that to give your inaccuracy. If $p = \frac{1}{2}$, every credence has the same expected inaccuracy from the vantage point of p ; if $p < \frac{1}{2}$, credence 0 minimizes expected inaccuracy from the vantage point of P ; and if $p > \frac{1}{2}$, credence 1 does so. So, while a is continuous, it is not strictly proper.³

A scoring rule measures the inaccuracy of a single credence. But the groups of individuals whose individual and group credences we will be assessing will typically have more than one credence; they will have credences in a range of different propositions. Adila, Benicio, and Cleo all have credences in H and in \bar{H} and no doubt also in many more propositions besides these. How are we to assess the accuracy, and thus epistemic value, of a credal state composed of more than one credence? We will adopt the natural answer. We will take the inaccuracy of a credence function — which of course represents a credal state consisting of credences in a number of propositions — to be the sum of the inaccuracies of the individual credences it assigns. That is, given a scoring rule s , we can define an inaccuracy measure \mathcal{J} as follows: \mathcal{J} takes a credence function cr and a possible state of the world w , and it returns

²Note that it is here that we appeal to the assumption that a credence of p in X is the *unique* rational response to your evidence. If your evidence merely made that credence one of a number of rational responses, we could not establish Strict Propriety. That would leave open the possibility that another credence $q \neq p$ in X is also a rational response to your evidence. And that would mean that we could not rule out the possibility that $ps(1, q) + (1 - p)s(0, q) \leq ps(1, p) + (1 - p)s(0, p)$.

³There are also strictly proper scoring rules that are not continuous (Schervish et al., 2009).

the inaccuracy of that whole credence function at w ; it takes the inaccuracy of cr to be the sum of the inaccuracies of the individual credences that cr assigns. In symbols:

$$\mathfrak{I}(cr, w) = \sum_{X \in \mathcal{F}} \mathfrak{s}_X(w(X), cr(X))$$

Here, \mathcal{F} is the set of propositions on which cr is defined; and $w(X)$ gives the truth value of proposition X in state of the world w — so $w(X) = 0$ if X is false at w , and $w(X) = 1$ if X is true at w . Thus, for instance, if we take the quadratic scoring rule \mathfrak{q} to be our scoring rule, it would generate this inaccuracy measure:

$$\mathfrak{I}_{\mathfrak{q}}(cr, w) = \sum_{X \in \mathcal{F}} \mathfrak{q}_X(w(X), cr(X)) = \sum_{X \in \mathcal{F}} |cr(X) - w(X)|^2$$

This, then, gives us our measure of epistemic *disvalue*. Pick a continuous and strictly proper scoring rule \mathfrak{s} ; generate from it an inaccuracy measure \mathfrak{I} in the way outlined above. $\mathfrak{I}(cr, w)$ is then the inaccuracy — and thus epistemic disvalue — of the credence function cr when the world is w . Given this, we can define the *expected* epistemic disvalue of one credence function from the point of view of another. Suppose cr is one credence function, and cr' is another. Then we can define the expected inaccuracy of cr' from the vantage point of cr as follows:

$$\text{Exp}_{\mathfrak{I}}(cr'|cr) = \sum_{w \in \mathcal{W}} cr(w) \mathfrak{I}(cr', w)$$

When this value is low, cr judges cr' to be doing well; when this value is high, cr judges cr' to be doing badly. Since \mathfrak{s} is strictly proper, and since \mathfrak{I} is generated from \mathfrak{s} in the way laid out above, it follows that there is an analogous sense in which \mathfrak{I} is also strictly proper: every credence function cr expects itself to be doing best. That is, for any two credence functions $cr \neq cr'$, $\text{Exp}_{\mathfrak{I}}(cr|cr) < \text{Exp}_{\mathfrak{I}}(cr'|cr)$. For this reason, when an inaccuracy measure \mathfrak{I} is generated from a continuous and strictly proper scoring rule \mathfrak{s} in the way laid out above, we say that \mathfrak{I} is an *additive and continuous strictly proper inaccuracy measure*.

4 The Accuracy Argument for Linear Pooling

I am now ready to make fully explicit my argument for linear pooling. It has three premises. Here is the first:

- (P1) The epistemic disvalue of a group credence function cr at a world w is given by $\mathfrak{I}(cr, w)$, where \mathfrak{I} is an additive and continuous strictly proper inaccuracy measure.

Here's the first thing this premise says: just as we can evaluate the epistemic value of the credence functions of individuals, so we can evaluate the epistemic value of the credence functions of groups. What's more, just as the epistemic value of an individual's credence function is its accuracy, so the epistemic value of a group credence function is its accuracy. This is plausible, but not uncontroversial. In many cases, we wish to identify a group credence function in order to use it to make decisions on behalf of the group. As a policymaker, I might ask what the scientific community thinks about future sea level rises because I want to decide ahead of time which mitigation strategies I should deploy. But we might also want to use it to say how the group, as a collective, represents the world. Plausibly, that's what

we're interested in when we ask about the authorship of *Hamlet*. We want to know something about how the scholarly community represents part of the world. Does that community think it most likely that Shakespeare wrote *Hamlet*, or that Edward de Vere did? Now, we know that individual credences play a dual role: they guide actions; but they also encode a representation of the world. Individual credences encode an individual's representation of the world; and that's why they have epistemic value, and why their epistemic value is their accuracy. The first thing that this first premise says is that group credences may also have this dual role: they guide the actions of that group; but they also encode facts about how that group represents the world. And, because of this, they have epistemic value as well as pragmatic value; and, moreover, their epistemic value is accuracy. The second thing that this premise says: the accuracy of group credence functions should be measured in the same way as the accuracy of individual credences; that is, by using an additive and continuous strictly proper inaccuracy measure. Now, it is worth noting that we will assume that, while individuals may disagree on the inaccuracy measure they use to give their own inaccuracy and the inaccuracy of other individuals, they all agree on the measure they use to give the inaccuracy of the group credence function. That is, while they might differ on the way in which they value inaccuracy for themselves and others like them, when it comes to the inaccuracy of a group credence, they agree on how to value its accuracy; and that agreement is recorded in a single additive and continuous strictly proper inaccuracy measure.

The second premise of the argument is this:

(P2) If, by the lights of every individual in a group, the expected epistemic value of one credence function is higher than the expected epistemic value of another credence function, then the latter cannot be the credence function of that group.

This premise belongs to a familiar family of claims about group attitudes. We might call these *unanimity preservation principles*. They all say that, while the individuals in the group will doubtless disagree on many things, when they agree about a certain sort of judgment and speak with one voice about it, their stance on that judgement should be respected in our account of their group attitude.

Recall our example from the start of the paper. Adila had credence 0.2 in H ; Benicio had 0.4; and Cleo had 0.5. Suppose I had then presented an account of group credences on which the group credence of Adila-Benicio-Cleo is 0.1. You would likely have found this strange. Part of this reaction, I think, is fuelled by a unanimity preservation principle: since all individuals agree that H is at least 20% likely and at most 50% likely, the group credence should also agree with this — a group credence of 0.1 in H does not.

But such principles must be handled with care. For instance, suppose we were to take the group credence in H to be 0.21. There doesn't seem to be anything troubling about that. And yet Adila, Benicio, and Cleo all agree that H is not 21% likely. This refutes the most general unanimity preservation principle, which says that, if each individual's credence function has a particular property, then the group credence function should also have that property. From this principle, it would follow that the group credence function must be the credence function of one of the individuals in the group — and that seems an extremely implausible demand.

So, if we accept (P2), it cannot be because we think it is *always* the case that, when the individuals speak with one voice, the group that they form should speak with that voice too. Nonetheless, (P2) seems very plausible. The group credence function is something put forward by the group as a summary of their diverse views. And it is answerable only to

them. Thus, if they all judge one credence function to be better than another, the latter could not possibly be the credence function they put forward as their collective view. And that is what (P2) says.

Note also that (P2) is a very natural application, in the epistemic case, of a principle that is very plausible in the practical case. Suppose that there are two mitigation strategies we might deploy against rising sea levels. And suppose that Adila, Benicio, and Cleo share exactly the same utility function, even though they differ in their credence function. Suppose that, from the vantage point of Adila's credence function, the expected utility of the first mitigation strategy is greater than the expected utility of the second. Suppose further that this is true from the vantage point of Benicio's credence function as well, and also from the vantage point of Cleo's credence function. That is, all three assign a higher expected utility to the first than to the second. That is, each of the three would, as good expected utility maximisers, choose the first over the second given the choice. Then it would seem bizarre to say that the group prefers the second to the first! And it would seem bizarre to say that the group endorses deploying the second strategy instead of the first. (P2) says the same in the epistemic case: if all three scientists agree on their ranking of two credence functions, then the group's ranking should agree; and the group credence function should not be outranked by some other credence function in the group's ranking.

Now, it is worth noting that, if we were to represent our agents' doxastic states as full beliefs, rather than as credences, the unanimity preservation principle just sketched would be more controversial. In that case, because of the Discursive Dilemma and related problems, it is difficult to give an account of group beliefs and group preferences that respects such a unanimity preservation principle (Pettit, 2007, 197-8).⁴ For instance, suppose we form group beliefs by a process of majority rule: that is, a group believes a proposition just in case a majority of its members believe that proposition. And suppose that we have a trio of policymakers, *A*, *B*, and *C*, all of whom will prefer option *a* to option *b* iff they believe each of the propositions *X*, *Y*, and *Z*; otherwise, they will prefer *b* to *a*. And suppose that *A* believes *X*, *Y*, but not *Z*; *B* believes *X*, *Z*, but not *Y*; and *C* believes *Y*, *Z*, but not *X*. So the group believes *X*, *Y*, and *Z*, since each proposition is believed by a majority of *A*, *B*, and *C*. Then each member prefers *b* to *a*, since each fails to believe one of *X*, *Y*, or *Z*. But the group prefers *a* to *b*, since it believes *X*, *Y*, and *Z*.

These sorts of problems don't arise in the credal case. Suppose each member of a group has the same utility function, and the group therefore has that same utility function. And suppose that the group credence function is obtained from the individual credence functions in accordance with Linear Pooling. Then if each member of the group assigns higher expected utility to option *a* than to option *b*, then so does the group — since the group credences are a weighted average of the individual credences and the group utility function is the utility function shared by all members of the group, the group's expected utilities are a weighted average of the individuals' expected utilities. As a result, there is no Discursive Dilemma in the credal case, and thus no threat to the sort of unanimity preservation principle that underpins (P2).

The third premise of our argument, (P3), is a mathematical theorem that acts as a bridge between (P1) and (P2), on the one hand, and Linear Pooling, on the other:

Theorem 1 *Suppose \mathcal{I} is an additive and continuous strictly proper inaccuracy measure. And suppose cr_1, \dots, cr_n is a collection of credence functions defined on the same set of propositions.*

⁴Thanks to Matt Kopec for urging me to address this concern.

(I) If Cr is not a linear pool of cr_1, \dots, cr_n , then there is Cr^* that is a linear pool of cr_1, \dots, cr_n such that, for each cr_i ,

$$\text{Exp}_{\mathfrak{J}}(Cr^*|cr_i) < \text{Exp}_{\mathfrak{J}}(Cr|cr_i)$$

(II) If Cr is a linear pool of cr_1, \dots, cr_n , then there is no $Cr^* \neq Cr$ such that, for each cr_i ,

$$\text{Exp}_{\mathfrak{J}}(Cr^*|cr_i) \leq \text{Exp}_{\mathfrak{J}}(Cr|cr_i)$$

We will not give the full proof here, though see the Appendix for a sketch of the proof.

Thus, we have the following argument for Linear Pooling:

(P1) The epistemic disvalue of a group credence function cr at a world w is given by $\mathfrak{J}(cr, w)$ (where \mathfrak{J} is an additive and continuous strictly proper inaccuracy measure).

(P2) If, by the lights of every individual in a group, the expected epistemic disvalue of one credence function is lower than the expected epistemic disvalue of another credence function, then the latter cannot be the credence function of that group.

(P3) Theorem 1.

Therefore,

(C) Linear Pooling.

Before we move on, it is worth saying how the present argument relates to another approach to group credences that appeals to scoring rules and inaccuracy measures. This is the approach taken by Sarah Moss (2011). As in our framework, we have a group of individuals, each with their own credence function. In general, in contrast with our approach in (P1), Moss allows that each individual might use a different inaccuracy measure to give the inaccuracy of the group credence function; however, in her argument for Linear Pooling, she assumes that they all use the same one. Moss then seeks the group credence function that somehow effects a compromise between the credence functions of the individuals who comprise the group. As we have seen above, given an individual's credence function cr and an inaccuracy measure \mathfrak{J} , we can evaluate the expected inaccuracy of any credence function cr' by the lights of that individual's credence function cr — it is $\text{Exp}_{\mathfrak{J}}(cr'|cr)$. Moss suggests that we can also evaluate the expected inaccuracy of any credence function *by the lights of the group*. She suggests that we take the group's expected inaccuracy of some credence function cr' to be a weighted average of the expected inaccuracies of cr' by the lights of the various individual credence functions held by the members of the group, and with respect to their individual ways of measuring inaccuracy. That is, if the individual credence functions of the members of a group G are cr_1, \dots, cr_n , and their inaccuracy measures are $\mathfrak{J}_1, \dots, \mathfrak{J}_n$, respectively, then there are weightings $\alpha_1, \dots, \alpha_n$ such that the group expected inaccuracy of a credence function cr is calculated as follows:

$$\text{Exp}_G(cr) = \sum_{i=1}^n \alpha_i \text{Exp}_{\mathfrak{J}_i}(cr|cr_i)$$

Moss then claims that the group credence function should be the credence function that minimizes Exp_G . What's more, she proves that, in the case in which each individual has the same inaccuracy measure (i.e. $\mathfrak{J}_i = \mathfrak{J}_j$, for each i, j) the credence function that minimizes

Exp_G is $Cr_G(-) = \alpha_1 cr_1(-) + \dots + \alpha_n cr_n(-)$ (Moss, 2011, 10). That is, Moss' method recommends Linear Pooling. This is an extremely interesting result, but I think it provides a weaker argument for Linear Pooling than the argument from premises (P1), (P2), and (P3) given above.

The problem is that Moss' argument assumes that we obtain the group's expected inaccuracy for a credence function cr' by taking a weighted arithmetic average of the individuals' expected inaccuracies for cr' . But the very question at issue here is how the numerically-represented judgments of individuals should be aggregated to give the numerically-represented judgments of the group that they form, whether those numerically-represented judgments are credences or expected inaccuracies. Taking weighted arithmetic averages is one method of aggregation — it is the one endorsed by Linear Pooling. But, as we saw at the beginning, there are others. But if you take the group's expected inaccuracy for cr' to be obtained by these other methods of aggregation, the credence function that minimizes the group's expected inaccuracy will not typically respect Linear Pooling. So, Moss' argument begs the question against those who think that the aggregation of numerically-represented judgments should proceed otherwise than by taking weighted arithmetic averages of those judgments.

The argument that I have presented here does not beg the same question. I make no assumption about how we aggregate the individual expected inaccuracies of a particular credence function to give the group expected inaccuracy of that credence function. Rather, I appeal to a unanimity preservation principle that says that, if every member of the group prefers credence function Cr^* to Cr , then Cr is not the group credence function.

5 Objections to Linear Pooling: preserving independences

As we mentioned above, the second premise of our argument, (P2), is a sort of unanimity preservation principle. And, as we saw, in conjunction with a standard account of the epistemic disvalue of a credal state, it can be used to establish Linear Pooling. However, it is often objected that precisely what is wrong with Linear Pooling is that it violates a plausible unanimity preservation principle. Linear pooling, as is often observed, does not preserve unanimous probabilistic independence (Laddaga, 1977; Lehrer & Wagner, 1983; Wagner, 1984; Genest & Wagner, 1987; Dietrich & List, 2015; Russell et al., 2015). We say that two propositions, H and E , are *probabilistically independent relative to a credence function* cr if $c(H|E) = c(H)$; that is, if the probability of H does not change when we condition on E . Equivalently, H and E are probabilistically independent relative to cr if $c(HE) = c(H)c(E)$; that is, if the probability of H and E both occurring is the probability of H occurring weighted by the probability of E occurring. Now, as is often observed, if two propositions are probabilistically independent relative to two credence functions cr and cr' , it is most likely that they will not be probabilistically independent relative to a mixture of cr and cr' . The following theorem, which is in the background in (Laddaga, 1977; Lehrer & Wagner, 1983), establishes this:

Theorem 2 *Suppose cr_1, cr_2 are credence functions, and $Cr = \alpha cr_1 + (1 - \alpha)cr_2$ is a mixture of them (that is, $0 \leq \alpha \leq 1$). Suppose that H and E are propositions and further that they are probabilistically independent relative to cr_1 and cr_2 . If H and E are also probabilistically independent relative to Cr , then at least one of the following is true:*

- (i) $\alpha = 0$ or $\alpha = 1$. That is, Cr simply is one of cr_1 or cr_2 .

(ii) $cr_1(H) = cr_2(H)$. That is, cr_1 and cr_2 agree on H .

(iii) $cr_1(E) = cr_2(E)$. That is, cr_1 and cr_2 agree on E .

On the basis of this well-known result, it is often said that there is a sort of judgment such that linear pooling does not preserve unanimity on that sort of judgment. The kind of judgment in question is judgment of independence. According to this objection to linear pooling, an individual judges that two propositions are independent whenever those propositions are probabilistically independent relative to her credence function. Thus, if your credence function is cr , and if H and E are probabilistically independent relative to cr , then you judge H and E to be independent. So, since two propositions can be independent relative to each of two different credence functions, but dependent relative to each of the non-extremal mixtures of those credence functions, linear pooling does not preserve unanimous judgments of independence — Adila and Benicio may be unanimous in their judgment that E is independent of H , while at the same time nearly all linear pools of their credences judge otherwise.

It seems to me that the mistake in this objection lies in its account of judgments of independence. I will argue that it is simply not the case that I judge H and E to be independent just in case my credence in H remains unchanged when I condition on E : it is possible to judge that H and E are independent without satisfying this condition; and it is possible to satisfy this condition without judging them independent. Let's see how.

First, suppose I am about to toss a coin. I know that it is either biased heavily in favour of heads or heavily in favour of tails. Indeed, I know that the objective chance of heads on any given toss is either 10% or 90%. And I know that every toss is stochastically independent of every other toss: that is, I know that, for each toss of the coin, the objective chance of heads is unchanged when we condition on any information about other tosses. Suppose further that I think each of the two possible biases is equally likely. I assign each bias a credence of 0.5. Then my credence that the coin will land heads on its second toss should also be 0.5. However, if I consider my credence in that same proposition *under the supposition that the coin landed heads on its first toss*, it is different — it is not 0.5. If the coin lands heads on the first toss, that provides strong evidence that the coin is biased towards heads and not tails — if it is biased towards heads, the evidence that it landed heads on the first toss becomes much more likely than it would if the coin is biased towards tails. And, as my credence that the coin has bias 90% increases, so does my credence that the coin will land heads on the second toss. So, while I know that the tosses of the coin are stochastically independent, the outcome of the first and the second toss are not probabilistically independent relative to my credence function.⁵

⁵More precisely: There are two possible objective chance functions ch_1 and ch_2 . If we let H_i be the proposition that the coin will land heads on its i^{th} toss, then the following hold:

- $ch_1(H_i) = 0.1$ and $ch_2(H_i) = 0.9$, for all i ;
- $ch_1(H_i H_j) = ch_1(H_i)ch_1(H_j)$ and $ch_2(H_i H_j) = ch_2(H_i)ch_2(H_j)$.

And if we let C_{ch_i} be the proposition that ch_i is the objective chance function, then given that I know that either ch_1 or ch_2 is the objective chance function, I should assign credence 1 to the disjunction of C_{ch_1} and C_{ch_2} . That is, $cr(C_{ch_1} \vee C_{ch_2}) = 1$. Now, given that H_1 and H_2 are independent relative to ch_1 and ch_2 , it seems natural to say that I judge H_1 and H_2 to be independent: I know that they are; and I assign maximal credence to a proposition, $C_{ch_1} \vee C_{ch_2}$, that entails that they are. Now suppose I think it equally likely that the coin has the 0.1 bias or that it has the 0.9 bias. So $cr(C_{ch_1}) = 0.5 = cr(C_{ch_2})$. Then, by the Principal Principle, my credence in heads on the second toss should be 0.5, for it should be $cr(H_2) = cr(C_{ch_1})ch_1(H_2) + cr(C_{ch_2})ch_2(H_2) = (0.5 \times 0.1) + (0.5 \times 0.9) = 0.5$. But suppose now that I condition on H_1 , the proposition that the coin lands heads on the first toss. If I were to learn H_1 , that would give me strong evidence that the coin is biased towards heads and not tails. After

Next, we can easily find examples in which two propositions are independent relative to my credence function, but I do not judge them independent. Indeed, there are examples in which I know for certain that they are not independent. Suppose, for instance, that there are just two probability functions that I consider possible chance functions. They agree on the chance they assign to HE and E , and thus they agree on the conditional chance of H given E . Both make H stochastically dependent on E . By the lights of the first, H depends positively on E — the conditional probability of H given E exceeds the unconditional probability of H . By the lights of the second, H depends negatively on E — the unconditional probability of H exceeds the conditional probability of H given E ; and indeed it does so by the same amount that the conditional probability of H given E exceeds the probability of H relative to the first possible chance function. Suppose I have equal credence in each of these possible chance hypotheses. Then my credence in H lies halfway between the chances of H assigned by the two possible chance functions. But, by hypothesis, that halfway point is just the conditional chance of H given E , on which they both agree. So my conditional credence in H given E is just my unconditional credence in H . So H and E are probabilistically independent relative to my credence function. Yet clearly I do not judge them stochastically independent. Indeed, I know them to be stochastically dependent — what I don't know is whether the dependence is positive or negative.⁶

So it seems that, whatever is encoded by the facts that make H and E probabilistically independent relative to my credence function, it is not my judgment that those two propositions are stochastically independent: I can know that H and E are stochastically independent without my credence function rendering them probabilistically independent; and I can know that H and E are stochastically dependent while my credence function renders them probabilistically independent. Perhaps, then, there is some other sort of independence that we judge to hold of H and E whenever our credence function renders those two propositions probabilistically independent? Perhaps, for instance, such a fact about our credence func-

all, the second chance hypothesis, C_{ch_2} , makes heads much more likely than does the first chance hypothesis, C_{ch_1} . And, indeed, again by the Principal Principle, $cr(H_2|H_1) = \frac{cr(H_2H_1)}{cr(H_1)} = \frac{cr(C_{ch_1})ch_1(H_2H_1) + cr(C_{ch_2})ch_2(H_2H_1)}{cr(C_{ch_1})ch_1(H_2) + cr(C_{ch_2})ch_2(H_2)} = \frac{(0.5 \times 0.1^2) + (0.5 \times 0.9^2)}{(0.5 \times 0.1) + (0.5 \times 0.9)} = \frac{0.82}{1} = 0.82 > 0.5 = cr(H_2)$. So, while I know that H_1 and H_2 are independent, and judge them so, it does not follow that they are independent relative to my credence function. The upshot: an individual might judge two propositions independent without those two events being probabilistically independent relative to her credence function.

⁶More precisely, suppose:

- (i) $ch_1(HE) = ch_2(HE)$ and $ch_1(E) = ch_2(E)$
- (ii) $ch_1(H|E) - ch_1(H) = ch_2(H) - ch_2(H|E) > 0$
- (iii) $cr(C_{ch_1}) = \frac{1}{2} = cr(C_{ch_2})$

First, note that:

$$cr(H|E) = \frac{cr(HE)}{cr(E)} = \frac{\frac{1}{2}ch_1(HE) + \frac{1}{2}ch_2(HE)}{\frac{1}{2}ch_1(E) + \frac{1}{2}ch_2(E)} = \frac{ch_i(HE)}{ch_i(E)} = ch_i(H|E)$$

Next, if we let $\beta = ch_1(H|E) - ch_1(H) = ch_2(H) - ch_2(H|E)$, then

$$\begin{aligned} cr(H) &= \frac{1}{2}ch_1(H) + \frac{1}{2}ch_2(H) \\ &= \frac{1}{2}(ch_1(H|E) - \beta) + \frac{1}{2}(ch_2(H|E) + \beta) \\ &= ch_i(H|E) = cr(H|E) \end{aligned}$$

tion encodes our judgment that H and E are *evidentially independent* or *evidentially irrelevant*? I think not. If you think that there are facts of the matter about evidential relevance, then these are presumably facts about which an individual may be uncertain. But then we are in the same position as we are with stochastic independence. We might have an individual who is uncertain which of two probability functions encodes the facts about evidential relevance. Each of them might make E epistemically relevant to H ; but it might be that, because of that individual's credences in the two possibilities, her credence function renders H and E probabilistically independent. If, on the other hand, you do not think there are facts of the matter about evidential relevance, it isn't clear how facts about my credence function could encode judgments about evidential relevance; nor, if they could, why we should care to preserve those judgments, even when they are made unanimously. Remember: we learned in section 4 that there will always be some features shared by all members of a group that cannot be shared with the group credence function.

Elkin & Wheeler (2016) try to dramatise the objection we are considering by presenting a Dutch Book argument against groups whose group credences fail to preserve independences shared by all members of the group. Their idea is this: Suppose that, relative to the credence function of each member of a group, propositions H and E are probabilistically independent. And suppose that, relative to their group credence function Cr , H and E are not probabilistically independent — that is, $Cr(HE) \neq Cr(H)Cr(E)$. Then, according to Elkin and Wheeler, there are two ways in which we can calculate the price at which the group will be prepared to buy or sell a £1 bet on the proposition HE — that is, a bet that pays £1 if HE turns out to be true, and which pays £0 if HE is false. First, the group will be prepared to buy or sell a £1 bet on HE at $£Cr(HE)$, since that is the group credence in HE . Second, Elkin and Wheeler claim that the group should also be prepared to buy or sell a £1 bet on HE at $£Cr(H)Cr(E)$, since $Cr(H)$ is the group credence in H , $Cr(E)$ is the group credence in E , and the group judges H and E to be independent. But, by hypothesis, $£Cr(HE) \neq £Cr(H)Cr(E)$, and if an agent has two different prices at which they are prepared to buy or sell bets on a given proposition, it is possible to Dutch Book them. Suppose that $Cr(H)Cr(E) < Cr(HE)$. Then we simply sell them a £1 bet on HE at $£Cr(HE)$, which they consider a fair price. This will give the group a net gain of $£(1 - Cr(HE))$ if HE is true and a net gain of $-£Cr(HE)$ if HE is false. And then we buy from them a £1 bet on HE at $£Cr(H)Cr(E)$, which is their other fair price. This will give the group a net gain of $£(Cr(H)Cr(E) - 1)$ if HE is true and a net gain of $£Cr(H)Cr(E)$ if HE is false. Thus, their total net gain if HE is true is $£(1 - Cr(HE)) + £(Cr(H)Cr(E) - 1) < £0$. And their total net gain if HE is false is $-£Cr(HE) + £Cr(H)Cr(E) < £0$. That is, the group is vulnerable to a series of bets, each of which it considers fair, but which collectively guarantee that it will lose money. And similarly with the signs reversed if $Cr(HE) < Cr(H)Cr(E)$.

The problem with this argument is the same as the problem with the original objection. The fact that H and E are probabilistically independent relative to each individual's credence function does not entail that each individual judges H and E to be independent. And without that, we have no reason to think that the group should also judge H and E independent, and thus no reason to think that the group should judge $£Cr(H)Cr(E)$ a fair price for a £1 bet on HE . In sum: I conclude that it does not count against linear pooling that it does not preserve probabilistic independence.

6 Objections to Linear Pooling: updating on evidence

Our second objection to linear pooling is closely related to the first. Both concern conditional probabilities — the first concerned their relationship to judgments of independence; the second concerns their relationship to rules of updating. As is often pointed out, linear pooling does not commute with updating by Bayesian conditionalization (Madansky, 1964; Genest, 1984; Dietrich & List, 2015; Berntson & Isaacs, 2013; Russell et al., 2015). The idea is this: Suppose that Adila and Benicio have credences in a range of propositions; and we take their group credence to be the linear pool of those credences determined by the weighting α for Adila and $1 - \alpha$ for Benicio. At this point, some new evidence arrives that is available to both members of the group. It comes in the form of a proposition that they both learn with certainty — perhaps they both learn the output from some climatological instrument. Bayesian conditionalization says that each individual, upon learning this evidence, should update their credences so that their new unconditional credence in a given proposition is just their old conditional credence in that proposition given the piece of evidence. How are we to update group credences in response to such evidence? There are two ways we might proceed: we might look to the individuals first, update their prior credence functions in accordance with the dictates of Bayesian conditionalization, and then take a linear pool of the resulting updated credence functions; or we might look to the group first, and update the group credence function in accordance with Bayesian conditionalization. Now suppose that, in the first approach, the weights used to pool the individual's posterior updated credence functions to give the group's posterior updated credence function are the same as the weights used to pool the individual's prior credence functions to give the group's prior credence function — that is, Adila's updated credence function is given weight α and Benicio's is given $1 - \alpha$. Then, in that situation, the two methods will rarely give the same result: updating and then linear pooling will most likely give a different result from linear pooling and then updating; or, as it is often put, linear pooling and updating do not commute. The following theorem makes this precise:

Theorem 3 ((Madansky, 1964)) *Suppose cr_1, cr_2 are credence functions, and $Cr(-) = \alpha cr_1(-) + (1 - \alpha)cr_2(-)$ is a mixture of them (that is, $0 \leq \alpha \leq 1$). And suppose that*

$$\alpha cr_1(H|E) + (1 - \alpha)cr_2(H|E) = Cr(H|E) \left(= \frac{\alpha cr_1(HE) + (1 - \alpha)cr_2(HE)}{\alpha cr_1(E) + (1 - \alpha)cr_2(E)} \right)$$

Then at least one of the following is true:

- (i) $\alpha = 0$ or $\alpha = 1$. That is, Cr simply is one of cr_1 or cr_2 .
- (ii) $cr_1(H|E) = cr_2(H|E)$. That is, cr_1 and cr_2 agree on H given E .
- (iii) $cr_1(E) = cr_2(E)$. That is, cr_1 and cr_2 agree on E .

This raises a problem for linear pooling, for it shows that the following are usually incompatible:

- (1) The rational update rule for individual credences is Bayesian conditionalization.
- (2) The rational update rule for group credences is Bayesian conditionalization.

- (3) Group credences are always obtained from individual credences in accordance with Linear Pooling.
- (4) The weights assigned to individuals do not change when those individuals receive a new piece of evidence.

The accuracy argument of section 4 seeks to establish (3), so we will not question that. And there are accuracy-based arguments in favour of Bayesian conditionalization as well. These are usually presented as establishing that *individuals* who plan to update otherwise than by conditionalization are irrational. For instance, Hilary Greaves & Wallace (2006) show that Bayesian conditionalization is the updating plan that, by the lights of the individual's prior credence function, uniquely minimizes the expected inaccuracy of the various possible posterior credence functions to which that updating plan might give rise, depending on the evidence it receives as input. And Briggs & Pettigrew (ms) show that, if an individual plans to update in any way other than by conditionalizing, there is an alternative plan they might have had that is guaranteed to result in less total inaccuracy for her when the inaccuracy of her initial credence function is added to the inaccuracy of her updated credence function. Surely, though, such arguments work just as well for the group credence functions. The argument of section 4 assumes that group credence functions, like individual credence functions, aim at accuracy. If that is so, then these two arguments for Bayesian conditionalization in the individual case tell just as strongly in favour of that updating rule in the group case. So we have (1) and (2).⁷

That leaves (4). Considerations of accuracy lead us to accept (1), (2), and (3); by doing so, they lead us to deny (4). In fact, this seems exactly right to me. To see why, let's begin by noting exactly how the weights must change to accommodate Bayesian conditionalization as the update plan for group credences in the presence of Linear Pooling. First, let's state the theorem, which is a particular case of the general result due to Howard Raiffa (1968, Chapter 8, Section 11):

Theorem 4 ((Raiffa, 1968)) *Suppose cr_1, cr_2 are credence functions, and $0 \leq \alpha, \alpha' \leq 1$. And suppose that*

$$\alpha' cr_1(H|E) + (1 - \alpha') cr_2(H|E) = \frac{\alpha cr_1(HE) + (1 - \alpha) cr_2(HE)}{\alpha cr_1(E) + (1 - \alpha) cr_2(E)}$$

Then at least one of the following is true:

(i)

$$\alpha' = \frac{\alpha cr_1(E)}{\alpha cr_1(E) + (1 - \alpha) cr_2(E)}$$

and

$$1 - \alpha' = \frac{(1 - \alpha) cr_2(E)}{\alpha cr_1(E) + (1 - \alpha) cr_2(E)}$$

⁷See (Leitgeb, 2016) for a view that accepts (3) and (4), but rejects (1) and (2). Leitgeb notes that there is an alternative updating rule, a certain sort of imaging, that does commute with linear pooling — indeed, it is the only one that does. This alternative updating rule is the extremal case of what Leitgeb & Pettigrew (2010) call *Alternative Jeffrey Conditionalization*, and which has since become known as *Leitgeb-Pettigrew* or *LP Conditionalization* (Levinstein, 2012). There is even an accuracy argument in its favour. Ben Levinstein (2012) raises worries about this updating rule; Richard Pettigrew (2016, Section 15.1) objects to the accuracy argument.

(ii) $cr_1(H|E) = cr_2(H|E)$. In this case, there are no restrictions on α' .

That is, to obtain the new weight, α' , for the first individual (whose initial credence function is cr_1), we take the old weight, α , we weight that by the credence that the first individual initially assigned to E , and we multiply by a normalizing factor. To obtain the new weight, $1 - \alpha'$, for the second individual (whose initial credence function is cr_2), we take the old weight, $1 - \alpha$, we weight that by the credence that the second individual initially assigned to E , and we multiply by the same normalizing factor. That is, the new weight that is assigned to an individual is determined entirely by her old weight and the initial credence she assigned to the proposition she has now learned to be true; her new weight is proportional to her old weight and the accuracy of her credence in that proposition. And indeed that seems exactly right. For we might think of these weights as encoding some facts about the expertise or reliability of the individuals in the group. Thus, when we learn a proposition, we increase the relative weighting of an individual in proportion to how confident they were in that proposition initially — that is, we reward their reliability with respect to this proposition by assigning them greater weight in the future.

Julia Staffel (2015, Section 6) objects to linear pooling on the grounds that it can only accommodate Bayesian conditionalization as the updating rule for individuals and groups by changing the weights assigned to the individuals in this way.⁸ Her worry is that, in certain cases, the required shifts in the weights are simply far more extreme than is warranted by the situation. Consider two polling experts, Nate and Ann. Over the course of their careers, they've been equally accurate in their predictions. As a result, when I ask for their group credence — the credence of Nate-Ann — I assign them equal weight: they both get a weight of 0.5. But then Ann has credence 0.8 in X and Nate has credence 0.2 in X , and X turns out to be true. When they both learn X , we have to shift the weights assigned to them in the group credence in order to preserve conditionalization — we have to shift Nate's from 0.5 to 0.2; and we have to shift Ann's from 0.5 to 0.8. That is, despite his long career of matching Ann's accuracy, one inaccurate prediction results in a drastic shift in the weight that Nate receives. Surely such an extreme shift is not justified by the situation. For instance, if Nate is now sceptical about a second proposition, Y , assigning it 0.1, while Ann is bullish, assigning it 0.9, then the group credence will be 0.74 — so, Nate's scepticism will do little to temper Ann's confidence.

I agree that such shifts are counterintuitive. However, I don't agree that this is a reason to reject Linear Pooling. Here's one reason not to worry about them: while a single more accurate credence in a proposition that is learned might increase Ann's weighting substantially, it can just as easily be reduced again if she has a single less accurate credence in some other truth that is learned later. Thus, if Nate and Ann next learn Y , a proposition to which Nate assigns 0.8 while Ann assigns 0.2, then the weighting of both immediately returns to 0.5. Thus, there need be no worry that a drastic shift in favour of one agent, once undertaken, will force the group credence function to stick close to that agent's credence function.

Another reason not to worry about drastic shifts: such shifts also occur in credences about chance hypotheses for any agent who satisfies the Principal Principle, a central tenet of Bayesian reasoning. Suppose I am in possession of a trick coin. You know that the bias of the coin towards heads is either 20% or 80%. You've watched 1,000 coin tosses: 500 came up heads; 500 tails. You began with credence 0.5 in each of the bias hypotheses. And you satisfy the Principal Principle at all times. This entails that, at each moment, your credence

⁸Thanks to Liam Kofi Bright, Julia Staffel, and Brian Weatherson for urging me to address this objection.

function is a linear pool of the possible chance functions, where the weight that you assign to a particular possible chance function is just your credence that it is the true chance function. As a result, having witnessed an equal number of heads and tails, your current credence in each of the bias hypotheses has returned to 0.5. But now you toss the coin again, and it lands heads. Then the Principal Principle and Bayesian conditionalization demand that your credence that the bias is 80% must shift to 0.8; and your credence in the bias is 20% must shift to 0.2. So, after a long run of equally good predictions, a single coin toss can shift your credences in the bias hypotheses dramatically. In fact, that single coin toss can shift your credences in the bias hypotheses exactly as dramatically as the weights assigned to individuals might shift if you adhere to Linear Pooling. And this is just a consequence of satisfying the innocuous and widely-accepted Principal Principle.⁹ This is my response to Staffel's objection.

Above, we presented a paradox: (1), (2), (3), and (4) are often incompatible. Which should we reject? In the light of the accuracy arguments just given for (1), (2), and (3), it seems to me that we should reject (4).

7 Conclusion

The individual credences of a group of individuals should be combined by linear pooling to give the collective credences of the group. If the group credences are not obtained in this way, there are alternative credences that every individual in the group expects to do better from an epistemic point of view; that is, each individual expects the alternative credences to be more accurate. This was the argument for Linear Pooling presented in section 4. In section 5, we noted that linear pools of individuals do not preserve the independences agreed upon by those individuals. But we concluded that this is no mark against Linear Pooling — these independences cannot be judgments of stochastic or evidential independence, and there is no other reason to preserve them. In section 6, we noted that linear pooling does not commute with conditionalization when the weight assigned to an individual remains unchanged after new evidence emerges. Again, we concluded that this does not tell against

⁹More precisely: There are two possible objective chance functions ch_1 and ch_2 . If we let H_i be the proposition that the coin will land heads on its i^{th} toss, then the following hold:

- $ch_1(H_i) = 0.2$ and $ch_2(H_i) = 0.8$, for all i ;
- $ch_1(H_i H_j) = ch_1(H_i)ch_1(H_j)$ and $ch_2(H_i H_j) = ch_2(H_i)ch_2(H_j)$

Let C_{ch_i} be the proposition that ch_i is the objective chance function. And let cr_i be my credence function after the i^{th} toss. Thus, by hypothesis, $cr_0(C_{ch_1}) = cr_0(C_{ch_2}) = 0.5$. Also, I assume that cr_i satisfies the Principal Principle at all times: that is,

$$cr_i(-|C_{ch_k}) = ch_k(-)$$

One consequence of this is:

$$cr_i(-) = cr_i(C_{ch_1})ch_1(-) + cr_i(C_{ch_2})ch_2(-)$$

Thus, my credence function at any point is a linear pool of the possible objective chance functions ch_1 and ch_2 , where the weights are determined by my credences in the chance hypotheses C_{ch_1} and C_{ch_2} . Now, after witnessing 500 heads and 500 tails, my credences are thus: $cr_{1,000}(C_{ch_1}) = cr_{1,000}(C_{ch_2}) = 0.5$. Now suppose I learn that the 1,001st toss landed heads — that is, I learn $H_{1,001}$. Then

$$cr_{1,001}(C_{ch_1}) = cr_{1,000}(C_{ch_1}|H_{1,001}) = cr_{1,000}(H_{1,001}|C_{ch_1}) \frac{cr_{1,000}(C_{ch_1})}{cr_{1,000}(H_{1,001})} = ch_1(H_{1,001}) \frac{0.5}{0.5} = ch_1(H_{1,001}) = 0.2$$

And similarly, $cr_{1,001}(C_{ch_2}) = 0.8$.

Linear Pooling — instead, it teaches us that weights should change in order to reward the accuracy of an individual’s prior credences and to punish their inaccuracy.

Linear pooling, then, places a requirement of rationality on the means by which the credences of individuals in a group are combined to give the collective credences of that group. The group credences of Adila, Benicio, and Cleo will be some mixture of their individual credences; the credence of the scholarly community that William Shakespeare wrote *Hamlet* must be some mixture of the credences assigned by the individual scholars that comprise that community; and so on.

A final note: The argument presented here has nothing to say about the interpretation of the weightings posited by Linear Pooling. Linear Pooling itself, but also this argument in its favour, doesn’t even claim that the most natural way for us to calculate or determine the group credence function is by determining these weights and then using them to produce the weighted average of the individual credence functions. All it claims is that, whatever the group credence function is, it had better be possible to recover it as a weighted average of the individual credence functions; there had better be some weights such that the weighted average given by those weights matches the group credence function. This does not require that the weights have any natural or plausible philosophical interpretation. What’s more, the argument presented here is agnostic about whether there are further rational restrictions on the weightings. It provides a necessary condition that must be satisfied by any candidate for the group credence function; but it does not pretend that it is also sufficient. There may be further restrictions. Whether there are or not, I leave to future work.

8 Appendix: sketch of proof of Theorem 1

The proof relies on three well-known facts, (BD1), (BD2a), (BD2b), below. All three concern a species of function called an *additive Bregman divergence*. For more detail on the mathematics of additive Bregman divergences, see (Banerjee et al., 2005), (Predd et al., 2009), (Pettigrew, 2016, 84-95).

- (BD1) Given an additive and continuous strictly proper inaccuracy measure \mathfrak{J} , there is an additive Bregman divergence \mathfrak{D} such that, for any two probabilistic credence functions cr, cr' ,

$$\mathfrak{D}(cr, cr') = \text{Exp}_{\mathfrak{J}}(cr'|cr) - \text{Exp}_{\mathfrak{J}}(cr|cr)$$

That is, the divergence from one credence function to another is given by subtracting the expected inaccuracy of the former by its own lights from the expected inaccuracy of the latter by the lights of the former.

- (BD2a) Given any additive Bregman divergence, if \mathcal{C} is a finite set of credence functions and cr' is outside the convex hull of \mathcal{C} , then there is cr^* in the convex hull of \mathcal{C} such that $\mathfrak{D}(cr, cr^*) < \mathfrak{D}(cr, cr')$, for all cr in \mathcal{C} .
- (BD2b) Given any additive Bregman divergence, if \mathcal{C} is a finite set of credence functions and cr' is inside the convex hull of \mathcal{C} , then, for any $cr^* \neq cr'$, we have $\mathfrak{D}(cr, cr') < \mathfrak{D}(cr, cr^*)$, for some cr in \mathcal{C} .

Proof of Theorem 1(I). Suppose Cr is not a linear pool of cr_1, \dots, cr_n . Then Cr lies outside the convex hull of the set of those credence functions. Thus, by (BD2a), there is Cr^* inside that convex hull such that $\mathfrak{D}(cr_i, Cr^*) < \mathfrak{D}(cr_i, Cr)$. And so, by (BD1), $\text{Exp}_{\mathfrak{J}}(Cr^*|cr_i) -$

$\text{Exp}_{\mathfrak{J}}(cr_i|cr_i) < \text{Exp}_{\mathfrak{J}}(Cr|cr_i) - \text{Exp}_{\mathfrak{J}}(cr_i|cr_i)$. And thus $\text{Exp}_{\mathfrak{J}}(Cr^*|cr_i) < \text{Exp}_{\mathfrak{J}}(Cr|cr_i)$, as required.

Proof of Theorem 1(II). Suppose Cr is a linear pool of cr_1, \dots, cr_n . Then Cr lies inside the convex hull of the set of those credence functions. Thus, by (BD2b), for any $Cr^* \neq Cr$, there is cr_i such that $\mathfrak{D}(cr_i, Cr) < \mathfrak{D}(cr_i, Cr^*)$. And so, by (BD1), $\text{Exp}_{\mathfrak{J}}(Cr|cr_i) - \text{Exp}_{\mathfrak{J}}(cr_i|cr_i) < \text{Exp}_{\mathfrak{J}}(Cr^*|cr_i) - \text{Exp}_{\mathfrak{J}}(cr_i|cr_i)$. And thus $\text{Exp}_{\mathfrak{J}}(Cr|cr_i) < \text{Exp}_{\mathfrak{J}}(Cr^*|cr_i)$, as required.

This completes the proof. □

References

- Aumann, R. J. (1976). Agreeing to Disagree. *The Annals of Statistics*, 4(6), 1236–1239.
- Banerjee, A., Guo, X., & Wang, H. (2005). On the Optimality of Conditional Expectation as a Bregman Predictor. *IEEE Transactions of Information Theory*, 51, 2664–69.
- Berntson, D., & Isaacs, Y. (2013). A New Prospect for Epistemic Aggregation. *Episteme*, 10(3), 269–281.
- Briggs, R., & Pettigrew, R. (ms). Conditionalization. Unpublished manuscript.
- Dietrich, F., & List, C. (2015). Probabilistic Opinion Pooling. In A. Hájek, & C. R. Hitchcock (Eds.) *Oxford Handbook of Philosophy and Probability*. Oxford: Oxford University Press.
- Elkin, L., & Wheeler, G. (2016). Resolving Peer Disagreements Through Imprecise Probabilities. *Nous*, doi: 10.1111/nous.12143.
- Genest, C. (1984). A characterization theorem for externally Bayesian groups. *Annals of Statistics*, 12(3), 1100–1105.
- Genest, C., & Wagner, C. (1987). Further evidence against independence preservation in expert judgement synthesis. *Aequationes Mathematicae*, 32(1), 74–86.
- Gibbard, A. (2008). Rational Credence and the Value of Truth. In T. Gendler, & J. Hawthorne (Eds.) *Oxford Studies in Epistemology*, vol. 2. Oxford University Press.
- Goldman, A. (2002). The Unity of the Epistemic Virtues. In *Pathways to Knowledge: Private and Public*. New York: Oxford University Press.
- Greaves, H., & Wallace, D. (2006). Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility. *Mind*, 115(459), 607–632.
- Heesen, R., & van der Kolk, P. (2016). A Game-Theoretic Approach to Peer Disagreement. *Erkenntnis*.
- Joyce, J. M. (1998). A Nonpragmatic Vindication of Probabilism. *Philosophy of Science*, 65(4), 575–603.
- Joyce, J. M. (2009). Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief. In F. Huber, & C. Schmidt-Petri (Eds.) *Degrees of Belief*. Springer.
- Laddaga, R. (1977). Lehrer and the consensus proposal. *Synthese*, 36, 473–77.
- Lehrer, K., & Wagner, C. (1983). Probability amalgamation and the independence issue: A reply to Laddaga. *Synthese*, 55(3), 339–346.
- Leitgeb, H. (2016). Imaging all the People. *Episteme*.

- Leitgeb, H., & Pettigrew, R. (2010). An Objective Justification of Bayesianism II: The Consequences of Minimizing Inaccuracy. *Philosophy of Science*, 77, 236–272.
- Levinstein, B. A. (2012). Leitgeb and Pettigrew on Accuracy and Updating. *Philosophy of Science*, 79(3), 413–424.
- Levinstein, B. A. (2015). With All Due Respect: The Macro-Epistemology of Disagreement. *Philosophers' Imprint*, 15(3), 1–20.
- Madansky, A. (1964). Externally Bayesian Groups. Memorandum rm-4141-pr, The RAND Corporation.
- Moss, S. (2011). Scoring Rules and Epistemic Compromise. *Mind*, 120(480), 1053–1069.
- Pettigrew, R. (2016). *Accuracy and the Laws of Credence*. Oxford: Oxford University Press.
- Pettit, P. (2007). Responsibility Incorporated. *Ethics*, 117, 171–201.
- Predd, J., Seiringer, R., Lieb, E. H., Osherson, D., Poor, V., & Kulkarni, S. (2009). Probabilistic Coherence and Proper Scoring Rules. *IEEE Transactions of Information Theory*, 55(10), 4786–4792.
- Raiffa, H. (1968). *Decision Analysis: Introductory Lectures on Choices under Uncertainty*. Reading: Addison-Wesley.
- Russell, J. S., Hawthorne, J., & Buchak, L. (2015). Groupthink. *Philosophical Studies*, 172, 1287–1309.
- Schervish, M. J., Seidenfeld, T., & Kadane, J. B. (2009). Proper Scoring Rules, Dominated Forecasts, and Coherence. *Decision Analysis*, 6(4), 202–221.
- Staffel, J. (2015). Disagreement and Epistemic Utility-Based Compromise. *Journal of Philosophical Logic*.
- Wagner, C. (1984). Aggregating subjective probabilities: some limitative theorems. *Notre Dame Journal of Formal Logic*, 25(3), 233–240.