

Aggregating agents with opinions about different propositions*

Richard Pettigrew

August 1, 2022

Abstract

There are many reasons we might want to take the opinions of various individuals and pool them to give the opinions of the group they constitute. If all the individuals in the group have probabilistic opinions about the same propositions, there is a host of pooling functions we might deploy, such as linear or geometric pooling. However, there are also cases where different members of the group assign probabilities to different sets of propositions, which might overlap a lot, a little, or not at all. There are far fewer proposals for how to proceed in these cases, and those there are have undesirable features. I begin by considering four proposals and arguing that they don't work. Then I'll describe my own proposal, which is intended to cover the situation in which we want to pool the individual opinions in order to ascribe an opinion to the group considered as an agent in its own right.

1 Introduction

There are many reasons we might want to take the opinions of various individuals and pool them to give the opinions of the group they constitute. They might be demographic modellers, and we wish to summarise their views for policymakers. Or they might be ice sheet modellers and we wish to pool the probabilities they assign to various future sea level scenarios in order to include these in our global climate models (Bamber & Aspinall, 2013; Bamber et al., 2019). We might be producing a textbook on the epidemiology of respiratory viruses, and we wish to present something that we might legitimately call the view of the scientific community (French, 1987, 2011). Or we might be the lead author on a scientific paper with many co-authors and we wish to ensure that the conclusions presented in the paper are genuinely those of the entire group of authors (Bright et al., 2017; Dang, 2019). Outside science, the individuals whose opinions we wish to pool might be employees of a company or institution whose collective opinion we wish to assess in order to determine liability for some harm, such as the board members of tobacco, oil, or social media companies, or the senior management of a university or a police force (Lackey, 2020). Or they might be superforecasters, renowned for the accuracy of their previous predictions of future political

*Acknowledgments: I'd like to extend my heartfelt thanks to Sylvia Wenmackers and two referees for this journal, all of whom generously gave extremely constructive comments on earlier versions of this material. This current version is dramatically improved as a result.

or sporting events, and we wish to learn what they, as a group, think about the outcome of a forthcoming election or the next World Cup (Tetlock & Gardner, 2015). And so on.

If all the individuals in the group have probabilistic opinions about the same propositions, there is a host of pooling functions we might deploy. For instance, linear pooling takes the group's probability for a proposition to be the arithmetic mean of the probabilities that its members assign to that proposition. Or, to calculate the group's probabilities for the possible states of the world, geometric pooling takes, for each state, the geometric mean of the probabilities that its members assign to that state, and then normalizes the results to ensure the pooled probabilities for the possible states sum to one. And so on. Each of these methods has its own desirable and undesirable features, which have been explored extensively (Genest & Zidek, 1986; Dietrich & List, 2015).

However, there are also cases where different members of the group assign probabilities to different sets of propositions, and these sets might overlap a lot, a little, or not at all. Indeed, unless the probabilities are elicited by asking the same roster of questions to each individual in the group, this is the situation we are most likely to encounter in the wild. For instance, if we glean the probabilities that academic experts assign by looking at what they report in their scholarly publications, we will find that they do not all report probabilities in the same propositions. One climate scientist might assign a probability to sea levels rising by at least 60cm by 2100, but nothing more fine-grained, while another might assign probabilities to it rising by 60-80cm, 80-100cm, and more than 100cm by that date. As they are usually formulated, most pooling functions don't cover these cases; more precisely, they don't tell us which credence the group assigns to a proposition to which some of its members fail to assign a credence. In this paper, I explore how we might fill that gap.

In Section 2, I'll introduce the formal framework in which we'll explore our problem. In Sections 3-6, I'll consider four proposals and argue that they don't work. Some of these exist in the literature explicitly as an answer to our question; some exist as answers to different questions, but are naturally repurposed to address ours; and some are simply occur to us naturally when we consider the question. Because so little has been written on this question, I begin with these four unsatisfactory proposals partly in order to clear the ground. But we will also see that, by doing so, an alternative proposal suggests itself. This is described in Section 8. It is designed to cover those situations in which our purpose in pooling the opinions of the individuals in the group is to assign an opinion to the group itself, considered as an agent in its own right.

2 The formal framework

Let me begin by laying out the formal framework we'll be working within.

- *Individuals* Let's assume that there are $n \geq 2$ individuals whose opinions we wish to pool.
- *Propositions* Let \mathcal{F}_i be the set of propositions to which individual i assigns subjective probabilities or degrees of belief, which we will call *credences* throughout. We might call \mathcal{F}_i their *agenda*. Let $\mathcal{F} = \bigcup_{i=1}^n \mathcal{F}_i$ be the union of all the individuals' agendas. Throughout, we assume that each \mathcal{F}_i is finite, and therefore \mathcal{F} is finite too.
- *Possible states of the world* Let \mathcal{W} be the set of possible worlds grained just finely enough to assign truth values to each proposition in \mathcal{F} . We might represent \mathcal{W} as the set of

classically consistent assignments of truth values to the propositions in \mathcal{F} . Since each \mathcal{F}_i is finite and therefore \mathcal{F} is finite, \mathcal{W} is also finite. If a proposition X in \mathcal{F} is true at world w in \mathcal{W} , we write $w \models X$, and we represent X by the set $\{w \in \mathcal{W} : w \models X\}$ of worlds at which it is true.

- *Subjective probabilities/credences* Let P_i record the credences that individual i assigns to the propositions in \mathcal{F}_i . We'll call this their *credence function*. For X in \mathcal{F}_i , $P_i(X)$ is the credence that individual i assigns to X . It is at least 0 and at most 1. We assume that these credences functions are *coherent*: that is, if \mathcal{F}_i^+ is the smallest Boolean algebra that includes \mathcal{F}_i , then it is possible to extend each P_i to a credence function P_i^+ on \mathcal{F}_i^+ that satisfies the probability axioms—that is, P_i^+ assigns credence 1 to the tautology, 0 to the contradiction, and the credence it assigns to a disjunction of pairwise incompatible propositions is the sum of the credences it assigns to the disjuncts.
- *Pooling functions* A pooling function Δ takes a sequence of n credence functions, P_1, \dots, P_n , where P_i assigns credences to the propositions in \mathcal{F}_i , and returns a credence function $\Delta(P_1, \dots, P_n)$, which assigns credences to the propositions in $\mathcal{F} = \bigcup_{i=1}^n \mathcal{F}_i$. In this definition, we don't assume that a pooling function must give a coherent output for any sequence of coherent inputs, but this is a desirable feature and in fact nearly all the examples we consider boast it.

Existing accounts of probabilistic opinion pooling deal with the particular case in which $\mathcal{F}_1 = \dots = \mathcal{F}_n = \mathcal{F}$. They often also assume that \mathcal{F} is a Boolean algebra.¹ In such cases, for every world w in \mathcal{W} , there is a proposition in \mathcal{F} that is true at w and only at w —these are sometimes called the *atoms* of the Boolean algebra \mathcal{F} . We abuse notation and write w for that proposition. We can then define linear and geometric pooling as follows:

Linear pooling Suppose P_1, \dots, P_n are defined on the same agenda \mathcal{F} . Then, if X is in \mathcal{F} , then

$$\Delta_{\text{LP}}(P_1, \dots, P_n)(X) = \frac{1}{n} \sum_{i=1}^n P_i(X)$$

That is, the credence that the linear pool of P_1, \dots, P_n assigns to a possible world is the arithmetic mean of the credences that each P_i assigns to it.

Geometric pooling Suppose P_1, \dots, P_n are defined on the same agenda \mathcal{F} , which is a Boolean algebra. And suppose that there is w in \mathcal{W} such that, for each P_i , $P_i(w) > 0$. Then, if w is in \mathcal{W} , then

$$\Delta_{\text{GP}}(P_1, \dots, P_n)(w) = \frac{\sqrt[n]{\prod_{i=1}^n P_i(w)}}{\sum_{w' \in \mathcal{W}} \sqrt[n]{\prod_{i=1}^n P_i(w')}}$$

And, for X in \mathcal{F} ,

$$\Delta_{\text{GP}}(P_1, \dots, P_n)(X) = \sum_{w \models X} \Delta_{\text{GP}}(P_1, \dots, P_n)(w)$$

That is, the credence that the geometric pool of P_1, \dots, P_n assigns to a possible world (or, more precisely, the corresponding atom of the algebra) is the normalized geometric mean of the credences that each P_i assigns to it; and the credence

¹For an exception, see (Dietrich & List, 2017).

it assigns to a proposition is the sum of the credences it assigns to the worlds at which the proposition is true (or, more precisely, the atoms that entail the proposition).

A couple of things to note:

- Since we assume throughout that each P_i is coherent, so is their linear pool and so is their geometric pool. In fact, we needn't even assume that each P_i is coherent in order to ensure that their geometric pool is coherent, but we do in order to ensure their linear pool is.
- Linear pooling is defined directly for each proposition in \mathcal{F} ; as a result, we need not assume anything about the structure of \mathcal{F} .
- Geometric pooling is defined first for the states of the world in \mathcal{W} , and then for each proposition in \mathcal{F} ; as a result, we must assume that \mathcal{F} contains the proposition $\{w\}$ for each w in \mathcal{W} .

In this paper, we ask: how should we pool in other cases? That is, how should we pool when two individuals have different agendas; that is, when $\mathcal{F}_i \neq \mathcal{F}_j$ for some individuals i and j ?

In the following four sections, I consider different answers to this question. None of them work. I consider them partly to situate my proposal within the literature and clear the ground, but also because solving the problem that rules out the first two proposals motivates the account that I will go on to give in the remainder of the paper. The third proposal also attempts to solve that problem. It fails for a different reason, but one that is equally illuminating. Those impatient to hear the solution I propose for a particular important case can skip to Section 8.

3 Extending linear and geometric pooling

As we saw in the previous section, linear and geometric pooling are only defined in the special case in which $\mathcal{F}_1 = \dots = \mathcal{F}_n = \mathcal{F}$; moreover, geometric pooling requires that \mathcal{F} is a Boolean algebra. But perhaps we might generalize them so that they apply when $\mathcal{F}_i \neq \mathcal{F}_j$ for some individuals i and j ?

For instance, suppose $\{X, Y, Z\}$ is a three-cell partition, and suppose the first of two individuals assigns credences to X , Y , and Z , so that $\mathcal{F}_1 = \{X, Y, Z\}$, while the second assigns credences only to X and Y , so that $\mathcal{F}_2 = \{X, Y\}$. Suppose their probability assignments are as follows:

	X	Y	Z
P_1	0.1	0.4	0.5
P_2	0.2	0.6	–

Then extending linear pooling to this case and taking the arithmetic means of the credences assigned to each gives:

	X	Y	Z
$\Delta_{LP}(P_1, P_2)$	0.15	0.5	0.5

But that's not coherent: the credences in X , Y , and Z sum to more than 1.

On the other hand, extending geometric pooling to this case and taking the geometric mean of the probabilities assigned to X , Y , and Z , and then normalizing, gives:

	X	Y	Z
$\Delta_{GP'}(P_1, P_2)$	0.125	0.433	0.442

Obviously that is coherent, because geometric pooling requires us to normalise the geometric means; so the result will always be coherent.

Perhaps we should follow the lead of geometric pooling and do this for our extended version of linear pooling in such cases as well? So first we take the arithmetic means, and then we normalise the result. That would give:

	X	Y	Z
$\Delta_{LP''}(P_1, P_2)$	0.1304	0.4347	0.4347

Unfortunately, both normalized extended linear pooling ($\Delta_{LP''}$) and extended geometric pooling ($\Delta_{GP'}$) violate a principle that I take to govern judgment pooling in the cases we are considering, where the agendas of some of our individuals differ.

Extension Invariance (EI) If, for each individual i , there is a unique coherent credence function P_i^* defined on $\mathcal{F} = \bigcup_{i=1}^n \mathcal{F}_i$ that extends P_i , then $\Delta(P_1, \dots, P_n) = \Delta(P_1^*, \dots, P_n^*)$.²

The point is well illustrated by the example we've been considering in this section. While P_2 does not assign a credence to Z , it does assign credences to X and Y and together those determine the credence it would have to assign to Z in order to remain coherent—since X, Y, Z form a partition, it must assign 0.2. Extension Invariance (EI) says that, in cases like this, where the probabilities that an individual assigns to the propositions in \mathcal{F}_i determine the probabilities they must assign to the remaining propositions in \mathcal{F} , the result of pooling the original probability assignments on $\mathcal{F}_1, \dots, \mathcal{F}_n$ should be the same as the result of pooling the probability functions on \mathcal{F} that are obtained by filling in the gaps in the way that coherence requires. The idea is that, if the credences you have reported commit you to further credences, then adding those further credences explicitly shouldn't change the outcome of pooling your credences with the credences of others. We will offer a partial accuracy-based justification of the principle in Section 7 below.

Thus, return to our case above:

	X	Y	Z
P_1	0.3	0.4	0.3
P_2	0.2	0.6	—
P_2^*	0.2	0.6	0.2

So (EI) says that $\Delta(P_1, P_2) = \Delta(P_1, P_2^*)$. But notice that neither normalized extended linear pooling ($\Delta_{LP''}$) nor extended geometric pooling ($\Delta_{GP'}$) deliver this:

	X	Y	Z
$\Delta_{LP''}(P_1, P_2)$	0.1304	0.4347	0.4347
$\Delta_{LP}(P_1, P_2^*)$	0.15	0.5	0.35
$\Delta_{GP'}(P_1, P_2)$	0.125	0.433	0.442
$\Delta_{GP}(P_1, P_2^*)$	0.149	0.517	0.333

² P_i^* defined on \mathcal{F} extends P_i defined on $\mathcal{F}_i \subseteq \mathcal{F}$ if $P_i^*(X) = P_i(X)$ for all X in \mathcal{F}_i . That is, if the restriction of P_i^* to \mathcal{F}_i is just P_i .

(EI) will cause problems for the proposal we consider in the following section as well. But before we move on to that, there is another problem with our attempt to extend linear and geometric pooling to the case in which $\mathcal{F}_i \neq \mathcal{F}_j$ for some i, j . Suppose $\mathcal{F}_1 = \{X \vee Y\}$ and $\mathcal{F}_2 = \{Y \vee Z\}$, where again X, Y , and Z form a partition. And suppose P_1 assigns credences only to the proposition in \mathcal{F}_1 , while P_2 assigns only to the proposition in \mathcal{F}_2 . In particular,

	$X \vee Y$	$Y \vee Z$
P_1	0.2	–
P_2	–	0.3

Now, first try to apply the extended linear pooling operator, $\Delta_{LP''}$. By averaging the credences in each proposition, we get:

	$X \vee Y$	$Y \vee Z$
$\Delta_{LP''}$	0.2	0.3

But that is incoherent: the credences in $X \vee Y$ and $Y \vee Z$ must sum at least to 1. So now we need to normalize. But how to do this? To normalize a credence function, we need to know the credences it assigns to the possible worlds. But in this case, we don't know that. So $\Delta_{LP''}(P_1, P_2)$ is undefined. And of course the same fate befalls $\Delta_{GP''}(P_1, P_2)$: indeed, it can't even get started, since it is defined initially on possible worlds, and then only at the second stage on logically weaker propositions.

3.1 A concern about Extension Invariance

You might think it is not reasonable to demand that our pooling function satisfy (EI). After all, it is easy to imagine cases in which, were an individual i to assign credences to all the propositions in \mathcal{F} , rather than merely those in \mathcal{F}_i , the credences they would assign to the propositions in \mathcal{F}_i would be different from the ones they actually assign. Here are two reasons this might happen. First, it might be the sort of case discussed in the literature on awareness growth, in which the individual becomes aware of a possibility they hadn't considered before and this leads them to reevaluate their opinions about the possibilities that they had considered before (Karni & Vierø, 2013; Wenmackers & Romeijn, 2016; Bradley, 2017; Steele & Stefánsson, 2021; Mahtani, 2021). For instance, I might assign credences only to the propositions *It will rain tomorrow* and *It will be sunny tomorrow*, and assign credence 50% to each, but then come to consider a third possibility, namely, *It will be misty tomorrow*, and that might lead me to reduce my credence in the original two propositions in order to assign some credence to this new one. Secondly, it might be a case in which, at the nearest world in which individual i has agenda \mathcal{F} rather than \mathcal{F}_i , their evidence is different. For instance, consider the example from the introduction in which one climate scientist assigns a credence only to the possibility that sea levels will rise by at least 60cm by 2100, but nothing more fine-grained, while another assigns credences to it rising by 60-80cm, 80-100cm, and more than 100cm by that date. Now, it might be that it is only climate scientists who work specifically on sea level modelling who assign credences to these more fine-grained possibilities. And it might be that such modellers have substantially different evidence from others. So, let's suppose that, at the actual world, the first climate scientist, who assigns credences only to the coarse-grained possibility, is not a sea level modeller. And now consider the nearest possible world in which they assign credences to the more fine-grained possibilities. In

that world, they are a sea level modeller and so their evidence is very different from what it is in the actual world. And that might lead them, in that world, to assign different credences to the coarse-grained possibility.

These situations are indeed possible. However, (EI) holds of the individuals in them all the same. After all, it is not justified by saying that, whenever there is, for each individual i , a unique coherent P_i^* defined on \mathcal{F} that extends P_i , this P_i^* gives the credences that the individual i would assign were there agenda \mathcal{F} instead of \mathcal{F}_i . It is not a counterfactual claim at all. Rather, as I sketched the justification above, (EI) is justified by noting that the credences that P_i assigns to the propositions in \mathcal{F}_i commit individual i to the credences that P_i^* assigns to the propositions in \mathcal{F} . So it does not say that a pooling function should give the same result whether applied to the individuals' actual credence functions or the credence functions they would have were they all to have \mathcal{F} as their agenda; it says that a pooling function should give the same result whether applied to the individuals' actual credence functions or to the credence functions on \mathcal{F} to which their actual credences commit them.

Notice that this justification for (EI) applies equally whether we use our pooling function to provide what Christian List (2014) calls *aggregate* or *corporate* group opinions. In List's terminology, an aggregate collective attitude provides a summary of the attitudes of the members of the collective, while a corporate collective attitude treats the group as an agent in its own right and ascribes to that agent the attitude in question. When we determine the sort of summary that is encoded in an aggregate group opinion, we surely wish to include not only the credences that the individuals have explicitly, but also those to which they are committed by those they have explicitly. And we surely do not wish to include the opinions they would have had in some nearby possible world in which they do explicitly assign credences to these other propositions. After all, we are summarising the group's actual opinions, not their counterfactual ones. And when we treat the group as an agent, we want to include in the supervenience basis for that group agent's opinions not only the credences the members have explicitly, but also those to which they are committed.

We will return to (EI) below. So far, we have appealed only to its intuitive plausibility. In Section 7, we will compare the accuracy of the credences you obtain if you use it with the accuracy of the credences you obtain if you violate it in various ways.

4 The Coherent Approximation Principle

In Section 3, we saw that it is difficult to extend linear and geometric pooling so that they apply to the problem of pooling credence functions defined on different agendas—that is, when $\mathcal{F}_i \neq \mathcal{F}_j$ for some i, j . In this section, we turn to one of the few treatments of the current problem from the literature. It is due to Daniel Osherson and Moshe Vardi (Osherson & Vardi, 2006).

In fact, Osherson and Vardi treat two problems at once. Not only do they not assume that the individuals to be pooled assign credences to the same propositions; they also do not assume that those individuals assign coherent credences. So they seek a pooling function that takes possibly incoherent credence functions over possibly different agendas and pools them into a coherent credence function on the union of the agendas. Their approach, which draws on the pioneering work of Sébastien Konieczny and Ramón Pino Pérez, is distance-based (Konieczny & Pino Pérez, 1998, 1999). That is, we begin by identifying a measure of distance from one credence to another. We then take the pool of a set of credence functions to

be the credence function for which the average of the sum of the distances from the credences that it assigns to the credences that the individuals assign is minimal. Osherson and Vardi consider two such measures of distance:

Absolute deviation For credences $0 \leq p, q \leq 1$,

$$AD(p, q) = |p - q|$$

Squared deviation For credences $0 \leq p, q \leq 1$,

$$SD(p, q) = |p - q|^2$$

And there are many others, including the popular Kullback-Leibler divergence:

Kullback-Leibler divergence For credences $0 \leq p \leq 1$ and $0 < q \leq 1$,

$$KL(p, q) = p \log \frac{p}{q} - p + q$$

We say that a measure \mathfrak{d} of distance from one credence to another is a *divergence* if (i) $\mathfrak{d}(p, q) \geq 0$ for all $0 \leq p, q \leq 1$ and (ii) $\mathfrak{d}(p, q) = 0$ iff $p = q$. AD, SD, and KL are all divergences. Now, given a divergence \mathfrak{d} , here is Osherson and Vardi's pooling function, where $\mathcal{P}_{\mathcal{F}}$ is the set of coherent credence functions on $\mathcal{F} = \bigcup_{i=1}^n \mathcal{F}_i$:

Coherent Approximation Principle $_{\mathfrak{d}}$ (CAP $_{\mathfrak{d}}$) For P_i defined on \mathcal{F}_i ,

$$\Delta_{\text{CAP}}^{\mathfrak{d}}(P_1, \dots, P_n) = \arg \inf_{P \in \mathcal{P}_{\mathcal{F}}} \sum_{i=1}^n \sum_{X \in \mathcal{F}_i} \mathfrak{d}(P(X), P_i(X))$$

That is, $\Delta_{\text{CAP}}^{\mathfrak{d}}(P_1, \dots, P_n)$ is the coherent credence function for which the average of the sums of the divergences from its credences to the credences assigned by P_1, \dots, P_n is minimal.³

In fact, if we wish the minimizer to be unique here, we must restrict the divergences that we use. For instance, recall our example from the previous section:

	X	Y	Z
P_1	0.1	0.4	0.5
P_2	0.2	0.6	–
P_2^*	0.2	0.6	0.2

Then, if we use the absolute deviation to measure the distance from one credence to another—that is, if $\mathfrak{d} = AD$ —then, providing $0.1 \leq P(X) \leq 0.2$, $0.4 \leq P(Y) \leq 0.6$, and $0.2 \leq P(Z) \leq 0.5$, P minimises the average distance to P_1 and P_2^* . Presumably for this reason, when Osherson writes about CAP again with different co-authors, they focus on squared deviation (Predd et al., 2008). We'll focus on squared deviation and Kullback-Leibler divergence for

³A more general definition would permit divergences between credence functions that are not generated by summing the divergences between the credences they assign. The results I present in this paper concern the less general definition, but this definition covers the vast majority of divergences that are actually considered in the literature. The more general definition must await future work. Thanks to an anonymous reviewer for this journal for this suggestion.

the moment.⁴ Here are the results of pooling P_1 and P_2 using $\Delta_{\text{CAP}}^{\text{SD}}$ and using $\Delta_{\text{CAP}}^{\text{KL}}$, and the results of pooling P_1 and P_2^* using $\Delta_{\text{CAP}}^{\text{SD}}$ and using $\Delta_{\text{CAP}}^{\text{KL}}$.

	X	Y	Z
$\Delta_{\text{CAP}}^{\text{SD}}(P_1, P_2)$	0.113	0.462	0.425
$\Delta_{\text{CAP}}^{\text{SD}}(P_1, P_2^*)$	0.15	0.5	0.35
$\Delta_{\text{CAP}}^{\text{KL}}(P_1, P_2)$	0.32	0.204	0.475
$\Delta_{\text{CAP}}^{\text{KL}}(P_1, P_2^*)$	0.148	0.389	0.463

Since the first and second row differ, $\Delta_{\text{CAP}}^{\text{SD}}$ violates (EI); since the third and fourth row differ, $\Delta_{\text{CAP}}^{\text{KL}}$ violates (EI).

Now, you might try to save the Coherent Approximation Principle in one of two ways. First, you might seek a divergence \mathfrak{d} for which $\Delta_{\text{CAP}}^{\mathfrak{d}}$ satisfies (EI). However, the following fact shows that this is impossible:

Proposition 1. *If \mathfrak{d} is differentiable in its first argument, $\Delta_{\text{CAP}}^{\mathfrak{d}}$ violates (EI).*

(The proof is given in the Appendix.)

Second, you might think that the problem arises because the single credence assigned to Z is given exactly as much weight as the two credences assigned to X and the two credences assigned to Y. But it's easy to check that assigning twice as much weight to $\mathfrak{d}(P(Z), P_1(Z))$ as to $\mathfrak{d}(P(X), P_1(X))$ or $\mathfrak{d}(P(Y), P_1(Y))$ doesn't bring the Coherent Approximation Principle into agreement with (EI). For instance,

$$(\text{SD}(P(X), 0.1) + \text{SD}(P(X), 0.2)) + (\text{SD}(P(Y), 0.4) + \text{SD}(P(Y), 0.6)) + 2 \times \text{SD}(P(Z), 0.5)$$

is minimized among coherent functions at $P = (0.1, 0.45, 0.45)$, while

$$\begin{aligned} &(\text{SD}(P(X), 0.1) + \text{SD}(P(X), 0.2)) + \\ &\quad (\text{SD}(P(Y), 0.4) + \text{SD}(P(Y), 0.6)) + \\ &\quad\quad (\text{SD}(P(Z), 0.5) + \text{SD}(P(Z), 0.2)) \end{aligned}$$

is minimized among coherent credence functions at $P = (0.15, 0.5, 0.35)$.

5 Pooling the sets of coherent credence functions that extend the individuals' credence functions

Like the Coherent Approximation Principle, the third proposal we'll consider asks us to pool by minimizing the average distance from some representation of the individuals' opinions. But whereas CAP represents individual i by the precise credences they implicitly assign to the propositions in \mathcal{F}_i , the third proposal represents them by the imprecise credences they assign to the propositions in \mathcal{F} . That is, instead of representing individual i by the *single credence function* P_i , which is defined on \mathcal{F}_i , we represent them by the following *set of credence functions*, which are defined on \mathcal{F} :

$$R_i = \{P \in \mathcal{P}_{\mathcal{F}} \mid (\forall X \in \mathcal{F}_i)[P(X) = P_i(X)]\}$$

⁴Pettigrew (2019) makes the same choice.

where $\mathcal{P}_{\mathcal{F}}$ is the set of coherent credence functions on \mathcal{F} , as above. So, R_i is the set of coherent extensions of P_i to \mathcal{F} . And we pool P_1, \dots, P_n by pooling R_1, \dots, R_n . And we pool R_1, \dots, R_n by finding a credence function P that minimizes the average distance from P to the R_i s. Now, there are two natural definition of the distance from P to R_i . On the first, it is the minimum distance between P and a member of R_i ; on the second, it is the maximum distance between P and a member of R_i . I'll consider both.

For many divergences and many P_1, \dots, P_n , these minimization problems will have a unique solution. In that case, we use the first definition of distance and define:

$$\Delta_{\text{MW}}^{\partial, \text{inf}}(P_1, \dots, P_n) = \arg \inf_{P \in \mathcal{P}_{\mathcal{F}}} \sum_{i=1}^n \left(\inf_{Q \in R_i} \sum_{X \in \mathcal{F}} \partial(P(X), Q(X)) \right)$$

And we use the second definition of distance and define:

$$\Delta_{\text{MW}}^{\partial, \text{sup}}(P_1, \dots, P_n) = \arg \inf_{P \in \mathcal{P}_{\mathcal{F}}} \sum_{i=1}^n \left(\sup_{Q \in R_i} \sum_{X \in \mathcal{F}} \partial(P(X), Q(X)) \right)$$

I use the subscript 'MW' for these pooling functions because this general method for combining sets of probability functions is proposed by Martin Adamčík and George Wilmers (Adamčík & Wilmers, 2014; Wilmers, 2015). Seamus Bradley (2019) criticizes it as a pooling function for sets of probability functions that represent uncertainty in the imprecise credence framework. But his criticisms are less worrying when it is used to pool sets of probability functions that represent gaps in credal reporting, as we do here, so I won't repeat them.

It is easy to see that these two pooling functions will satisfy (EI). After all, if there is a unique coherent credence function P_i^* , defined on \mathcal{F} , that extends P_i , which is defined on \mathcal{F}_i , then the set of coherent probability functions that extends P_i is the same as the set of coherent probability functions that extends P_i^* —both contain only P_i^* . That is:

$$R_i = \{P : \mathcal{F} \rightarrow [0, 1] \mid P \in \mathcal{P}_{\mathcal{F}} \ \& \ (\forall X \in \mathcal{F}_i)[P(X) = P_i(X)]\} = \\ \{P_i^*\} = \{P : \mathcal{F} \rightarrow [0, 1] \mid P \in \mathcal{P}_{\mathcal{F}} \ \& \ (\forall X \in \mathcal{F}_i)[P(X) = P_i^*(X)]\} = R_i^*$$

So these proposals do not suffer from the same problem as the previous two. But they do face a problem: they give implausible answers in reasonably straightforward cases. For instance, suppose $\mathcal{F} = \{X, Y, Z\}$, where X, Y , and Z form a partition, and $\mathcal{F}_1 = \{X\}$ and $\mathcal{F}_2 = \{Y\}$. And suppose

	X	Y	Z
P_1	0.8	–	–
P_2	–	0.8	–

So:

- $R_1 = \{P \in \mathcal{P}_{\mathcal{F}} : P(X) = 0.8\}$
- $R_2 = \{P \in \mathcal{P}_{\mathcal{F}} : P(Y) = 0.8\}$

We can illustrate these two sets of probabilities by plotting them within the three-dimensional simplex on a barycentric plot (see Figure 1). The problem is that, if ∂ is squared deviation




Figure 1: The barycentric plot of the 2-simplex with $(1, 0, 0)$ at bottom left, $(0, 1, 0)$ at bottom right, and $(0, 0, 1)$ at the top. The dotted lines represent R_1 and R_2 , respectively. And the result of applying $\Delta_{MW}^{\delta,inf}$ and $\Delta_{MW}^{\delta,sup}$ to P_1 and P_2 is plotted.

(SD) or Kullback-Leibler divergence (KL), then $\Delta_{CAP}^{\delta,inf}(P_1, P_2)$ and $\Delta_{CAP}^{\delta,sup}(P_1, P_2)$ are as follows:

	X	Y	Z
$\Delta_{MW}^{\delta,inf}(P_1, P_2)$	0.5	0.5	0
$\Delta_{MW}^{\delta,sup}(P_1, P_2)$	0.4	0.4	0.2

These are plotted on the simplex as well. The problem here is that both seem too extreme. $\Delta_{MW}^{\delta,inf}(P_1, P_2)$ assigns credence 0 to Z , even though nothing in the opinions of either agent forces that. It is the same pool we would obtain if both agents were to assign credence 0 to Z and fill in Y in such a way that they remained coherent. That is,

$$\Delta_{MW}^{\delta,inf}(P_1, P_2) = \Delta_{MW}^{\delta}(P_1^{\circ}, P_2^{\circ})$$

where

	X	Y	Z
P_1°	0.8	0.2	0
P_2°	0.2	0.8	0

And $\Delta_{MW}^{\delta,sup}(P_1, P_2)$ assigns credence 0.2 to Z , even though nothing in the opinions of either agent forces that. It is the same pool we would obtain if agent 1 were to assign credence 0 to Y and fill in Z in such a way that they remain coherent, and agent 2 were to assign credence 0 to X and fill in Z in such a way that they remain coherent. That is,

$$\Delta_{MW}^{\delta,sup}(P_1, P_2) = \Delta_{MW}^{\delta}(P_1^{\dagger}, P_2^{\dagger})$$

where

	X	Y	Z
P_1^\dagger	0.8	0	0.2
P_2^\dagger	0	0.8	0.2

6 Maximising entropy within the set of possible pools

Here's another proposal that arises naturally. Let

$$R_{LP} = \{\Delta_{LP}(P'_1, \dots, P'_n) : P'_1 \in R_1, \dots, P'_n \in R_n\}$$

That is, R_{LP} is the set of linear pools of coherent extensions of the individuals' credence functions. Then let the pool of P_1, \dots, P_n be the credence function in R_{LP} with maximum entropy.⁵ First, define the Shannon entropy of a probability function P defined over a set \mathcal{W} of possible worlds as follows (Shannon, 1948):

$$H(P) = - \sum_{w \in \mathcal{W}} P(w) \log P(w)$$

Then let

$$\Delta_{ME}^{LP}(P_1, \dots, P_n) := \arg \max_{P \in R_{LP}} H(P)$$

The problem with this approach is that it gives the same implausible answer as $\Delta_{MW}^{\delta, \sup}$ gave in the case we considered in the previous section. That is, if

	X	Y	Z
P_1	0.8	–	–
P_2	–	0.8	–

Then

$$\Delta_{ME}^{LP}(P_1, P_2) \mid \begin{array}{c|ccc} & X & Y & Z \\ \hline & 0.4 & 0.4 & 0.2 \end{array}$$

Again, we illustrate this in a barycentric plot—see Figure 2.

7 Extension Invariance and the accuracy of pooling functions

In Sections 3 and 4, we criticized the extensions of linear and geometric pooling, Δ_{LP}'' and Δ_{GP}' , and the Coherent Approximation Principle, Δ_{CAP}^0 , because they both violate (EI), the principle that says that, when there's a unique coherent extension of each credence function to the full algebra, pooling those extensions should give the same result as pooling the original credence functions. At that point, I merely appealed to the intuitive force of (EI); I gave no further argument in its favour. But there is something to be said for pooling functions that satisfy it, at least when they are compared with CAP.

⁵Strictly speaking, this will only work if $\mathcal{F} = \bigcup_{i=1}^n \mathcal{F}_i$ contains \mathcal{W} . It's an interesting question how the proposal might be extended beyond this, perhaps by considering the extensions of each P_i not only to \mathcal{F} but to \mathcal{F}^* , the smallest Boolean algebra that contains \mathcal{F} . But, as we will see, the proposal doesn't work, so I won't spend time on that.




Figure 2: The barycentric plot of the 2-simplex with $(1, 0, 0)$ at bottom left, $(0, 1, 0)$ at bottom right, and $(0, 0, 1)$ at the top. The dotted lines represent R_1 and R_2 , respectively. And the result of applying Δ_{ME}^{LP} to P_1 and P_2 is plotted.

Let's begin with a slight adaptation of the simple example from above:

	X	Y	Z
P_1	0.1	0.4	—
P_2	0.2	0.6	—
P_1^*	0.1	0.4	0.5
P_2^*	0.2	0.6	0.2

(EI) says that pooling P_1 and P_2 should give the same result as pooling P_1 and P_2^* , which should give the same result as pooling P_1^* and P_2 , which should give the same result as pooling P_1^* and P_2^* . That is, if Δ is our pooling function,

$$\Delta(P_1, P_2) = \Delta(P_1, P_2^*) = \Delta(P_1^*, P_2) = \Delta(P_1^*, P_2^*)$$

But let's apply CAP using the squared deviation:

	X	Y	Z
$\Delta_{CAP}^{SD}(P_1, P_2)$	$\frac{12}{80}$	$\frac{40}{80}$	$\frac{28}{80}$
$\Delta_{CAP}^{SD}(P_1, P_2^*)$	$\frac{15}{80}$	$\frac{43}{80}$	$\frac{22}{80}$
$\Delta_{CAP}^{SD}(P_1^*, P_2)$	$\frac{9}{80}$	$\frac{37}{80}$	$\frac{34}{80}$
$\Delta_{CAP}^{SD}(P_1^*, P_2^*)$	$\frac{12}{80}$	$\frac{40}{80}$	$\frac{28}{80}$

Now, notice that $\Delta_{\text{CAP}}^{\text{SD}}(P_1^*, P_2^*)$ is the same as $\Delta_{\text{CAP}}^{\text{SD}}(P_1, P_2)$, and both are the midpoint between $\Delta_{\text{CAP}}^{\text{SD}}(P_1, P_2^*)$ and $\Delta_{\text{CAP}}^{\text{SD}}(P_1^*, P_2)$. That is,

$$\begin{aligned}\Delta_{\text{CAP}}^{\text{SD}}(P_1^*, P_2^*) &= \Delta_{\text{CAP}}^{\text{SD}}(P_1, P_2) = \\ &= \frac{1}{2}(\Delta_{\text{CAP}}^{\text{SD}}(P_1, P_2^*) + \Delta_{\text{CAP}}^{\text{SD}}(P_1^*, P_2)) = \\ &= \frac{1}{4}(\Delta_{\text{CAP}}^{\text{SD}}(P_1^*, P_2^*) + \Delta_{\text{CAP}}^{\text{SD}}(P_1, P_2) + \Delta_{\text{CAP}}^{\text{SD}}(P_1, P_2^*) + \Delta_{\text{CAP}}^{\text{SD}}(P_1^*, P_2))\end{aligned}$$

That is, when we include the credal assignment to Z that P_1 determines, but not the assignment that P_2 determines, $\Delta_{\text{CAP}}^{\text{SD}}$ pulls the pool towards P_1 and away from P_2 ; and, mutatis mutandis, if we include the credal assignment to Z that P_2 determines, but not the one that P_1 determines. And, moreover, the pull is the same but in opposite directions in the two cases. So, when we average them, we obtain what we would have obtained if we'd left out both assignments to Z (and pooled P_1 and P_2) or if we'd included both assignments to Z (and pooled P_1^* and P_2^*).

What does this tell us? Well, suppose our favoured pooling function for those cases in which all individuals have the same agenda is linear pooling; and suppose we extend that pooling function in line with (EI). Then we can say that following in favour of our approach and against CAP. First, we note the following corollary of the Diversity Prediction Theorem (Galton, 1907; Page, 2007):

Theorem 2. *For any \mathcal{F} and credence functions Q, Q_1, \dots, Q_n defined on \mathcal{F} ,*

$$\sum_{X \in \mathcal{F}} \text{SD}(\Delta_{\text{LP}}(Q_1, \dots, Q_n)(X), Q(X)) < \frac{1}{n} \sum_{i=1}^n \sum_{X \in \mathcal{F}} \text{SD}(Q_i(X), Q(X))$$

This says that, for any credence function Q and any set of credence functions Q_1, \dots, Q_n all defined on the same set of propositions, the distance of the linear pool of Q_1, \dots, Q_n from Q is always less than the average distance of the Q_i s from Q , when the distance between credences is measured using squared deviation.⁶

How does this help? Well, given a possible world w , let V_w be the credence function that assigns maximal credence to all propositions that are true at w and minimal credence to all propositions that are false at w : that is, $V_w(X) = 1$ if X is true at w , and $V_w(X) = 0$ if X is false at w . We might call V_w the omniscient credence function. It is natural to say that the ideal credence function for an individual to have at a world is the omniscient credence function at that world, and that a credence function is more accurate the closer it lies to that omniscient credence function. So we might say that the inaccuracy of a credence function Q_i at world w is the sum of the squared deviations between the credences it assigns and the credences that V_w assigns: we call this the Brier score of inaccuracy. So, if P is defined on \mathcal{F} ,

$$\mathfrak{B}(P, w) = \sum_{X \in \mathcal{F}} (P(X) - V_w(X))^2$$

And we might think that a credence function is doing better, epistemically speaking, the greater its inaccuracy and the lower its Brier score. That is, P is better than Q at w just in case $\mathfrak{B}(P, w) < \mathfrak{B}(Q, w)$ (Brier, 1950; Rosenkrantz, 1981; Pettigrew, 2016). Then, by Theorem 2,

⁶Indeed, this generalizes to any convex divergence, but I'll focus on squared deviation here for the sake of concreteness.

Corollary 3. For any \mathcal{F} , any world w , and any credence functions Q_1, \dots, Q_n defined on \mathcal{F} ,

$$\mathfrak{B}(\Delta_{\text{LP}}(Q_1, \dots, Q_n), w) < \frac{1}{n} \sum_{i=1}^n \mathfrak{B}(Q_i, w)$$

That is, the inaccuracy of the linear pool of Q_1, \dots, Q_n is less than the average inaccuracy of the Q_i s.

Now, recall that the linear pool of $\Delta_{\text{CAP}}^{\text{SD}}(P_1^*, P_2^*)$, $\Delta_{\text{CAP}}^{\text{SD}}(P_1, P_2)$, $\Delta_{\text{CAP}}^{\text{SD}}(P_1, P_2^*)$, and $\Delta_{\text{CAP}}^{\text{SD}}(P_1^*, P_2)$ is just $\Delta_{\text{LP}}(P_1^*, P_2^*)$. Then it follows that, for any world, the inaccuracy of $\Delta_{\text{LP}}(P_1^*, P_2^*)$ at that world is less than the average inaccuracy of $\Delta_{\text{CAP}}^{\text{SD}}(P_1^*, P_2^*)$, $\Delta_{\text{CAP}}^{\text{SD}}(P_1, P_2)$, $\Delta_{\text{CAP}}^{\text{SD}}(P_1, P_2^*)$, and $\Delta_{\text{CAP}}^{\text{SD}}(P_1^*, P_2)$ at that world. That is,

$$\begin{aligned} \mathfrak{B}(\Delta_{\text{LP}}(P_1^*, P_2^*), w) < \\ \frac{1}{4} (\mathfrak{B}(\Delta_{\text{CAP}}^{\text{SD}}(P_1^*, P_2^*), w) + \mathfrak{B}(\Delta_{\text{CAP}}^{\text{SD}}(P_1, P_2), w) + \\ \mathfrak{B}(\Delta_{\text{CAP}}^{\text{SD}}(P_1, P_2^*), w) + \mathfrak{B}(\Delta_{\text{CAP}}^{\text{SD}}(P_1^*, P_2), w)) \end{aligned}$$

So pooling in line with (EI) is more accurate than pooling in line with CAP, at least in expectation and if you are equally likely to find yourself pooling P_1 and P_2 as you are to find yourself pooling P_1 and P_2^* , or P_1^* and P_2 , or P_1^* and P_2^* .

Does this generalise beyond the specific case of P_1 and P_2 ? Yes, as the following theorem shows:

Theorem 4. Suppose $\mathcal{F}, \mathcal{F}'$ are two sets of propositions and $\mathcal{F}' \subseteq \mathcal{F}$. Suppose P_1 is a credence function on \mathcal{F}' and P_1^* is the unique coherent extension of P_1 to \mathcal{F} ; and suppose P_2 is a credence function on \mathcal{F}' and P_2^* is the unique coherent extension of P_2 to \mathcal{F} . Then

$$\Delta_{\text{LP}}(P_1^*, P_2^*) = \frac{1}{4} (\Delta_{\text{CAP}}^{\text{SD}}(P_1, P_2) + \Delta_{\text{CAP}}^{\text{SD}}(P_1^*, P_2^*) + \Delta_{\text{CAP}}^{\text{SD}}(P_1^*, P_2) + \Delta_{\text{CAP}}^{\text{SD}}(P_1, P_2^*))$$

Now, suppose you enter a pooling task knowing only that the individuals will assign credences either to the propositions in \mathcal{F}' or to the propositions in \mathcal{F} , where $\mathcal{F}' \subseteq \mathcal{F}$. Then, if we assume that there is no correlation between the particular credences the individuals assign and whether they assign them only to the propositions in \mathcal{F}' or to the propositions in \mathcal{F} , then it is as likely that the group whose opinions you wish to pool consists of P_1 and P_2^* as it is that it will consist of P_1^* and P_2 , and as likely that it consists of P_1 and P_2 , and as likely that it consists of P_1^* and P_2^* . And if that's right then the expected inaccuracy of using a rule that respects (EI) and thus sets $\Delta(P_1^*, P_2) = \Delta(P_1, P_2^*) = \Delta_{\text{LP}}(P_1^*, P_2^*)$ is lower than the expected inaccuracy of using CAP, which will give each of $\Delta_{\text{CAP}}^{\text{SD}}(P_1, P_2^*)$, $\Delta_{\text{CAP}}^{\text{SD}}(P_1^*, P_2)$, $\Delta_{\text{CAP}}^{\text{SD}}(P_1, P_2)$, and $\Delta_{\text{CAP}}^{\text{SD}}(P_1^*, P_2^*)$ a probability of 25%.

8 Beyond Extension Invariance

Extension Invariance (EI) tells us how our pooling function should work when, for each individual i , there is a unique coherent extension P_i^* of P_i from \mathcal{F}_i to \mathcal{F} . In such a case, (EI) tells us, you pick the pooling function you favour for those cases in which all the credence functions to be pooled are defined on the same set of propositions, and you apply it to the extended credence functions P_1^*, \dots, P_n^* , which are all defined on \mathcal{F} . As it stands, however, (EI) does not tell us how to proceed when, for some individual i , there is more than one

coherent credence function that extends P_i from \mathcal{F}_i to \mathcal{F} . In this final section, I consider an important sort of case in which we face this problem and propose a solution for that case.

In the case I want to consider, we pool credences in order to give what Christian List (2014) calls a corporate collective attitude. Recall from above: in List's terminology, when we ascribe a corporate collective attitude, we assert first that the group counts as an agent in its own right and second that this group agent has the attitude in question. From the examples given in the introduction, these are the ones I envisage falling into this category: the epidemiologists of viruses whose views we wish to present as the view of the scientific community in our textbooks; the co-authors on a multi-authored scientific paper whose group view as a collective author we wish to present to the scientific community; and the employees of a company or institution whose collective view we wish to identify in order to assess liability for some harm. In all three cases, the group agents play a role in some normative enterprise. In the first two, it is the enterprise of science, which has norms that govern the assertions included in textbooks and scientific papers. In the third, it is the legal system, and there are norms here that govern the beliefs we ascribe to an individual whose liability for some harm we are assessing.

While these normative enterprises and the roles within them that the groups play are quite different, I will argue that a similar norm governs how we should pool the credences of the individuals in such a group to give the group's credences. It is a conservative norm. It says that we should first pick, for each individual i , a particular credence function P_i^* that extends P_i from \mathcal{F}_i to \mathcal{F} ; in particular, it says that we should pick P_i^* in the most conservative and unopinionated way possible; that is, we should introduce as little in the way of further opinions as we can when we extend; and then, second, we should aggregate these extended credence functions using whatever pooling function we favour for those cases in which all credence functions are defined on the same agenda—perhaps linear pooling, perhaps geometric pooling, perhaps something else.

Why is this the appropriate norm in the scientific case? In particular, why does the norm require us to extend P_i to P_i^* in the most conservative way possible. In fact, I think there are two reasons. The first reason is the duty of the textbook's author or the paper's lead author to represent fairly the views of the individuals on behalf of whom they write. The textbook's author presents the views of that part of the scientific community; the lead author presents the views of their fellow co-authors. In both cases, they have a duty not to impute to those individuals any further opinions beyond what is necessary to extend their credences to the full set \mathcal{F} . The second reason is the duty of scientific authors to their audience. Now, I don't think it is the duty of each scientist not to form opinions beyond what is strictly implied by their evidence. Over years of training and experience in their field, scientists gain an ability to form opinions on the basis of the evidence that sometimes seems to go beyond what the non-expert might conclude, and yet which it is legitimate to report in a scientific publication because of the expertise of the scientist. Nonetheless, when the scientist hasn't formed any opinion about a proposition and when we must nonetheless ascribe an opinion to them in order to carry out the pooling, we are obliged to make that opinion as conservative as possible. In other words, deviations from a sort of Cliffordian conservatism about opinion are permitted, but only when they are made explicitly by the scientist, and not when they are made by a textbook author or lead author on a paper who is filling in the gaps in another scientist's opinions.

Why is conservatism the appropriate norm in the legal case? Here, I think the key lies in the legal notion of the 'reasonable person'. Often this abstract individual is invoked to

personify a certain standard of proof that is required in order to find a defendant liable or guilty. On the websites of many US police departments, you will find a definition of ‘probable cause’ in terms of what a ‘reasonable person’ would believe on the basis of the evidence in hand. But it is also used to determine when a defendant’s actions are reasonable. For instance, in *Brown vs. Kendall*, Chief Justice Shaw determined that the ‘ordinary care’ that is necessary for the defendant to avoid liability is “that kind and degree of care, which prudent and cautious men would use”.⁷ And in *Commonwealth vs. Horsfall*, Chief Justice Rugg declared that “every traveller upon a highway is bound to exercise the care of the ordinarily prudent and cautious person under all circumstances”.⁸ In both of these cases, we see that the ‘reasonable person’ is identified with the ‘prudent and cautious person’. In the cases cited, the prudence and caution relate to the individual’s practical choices about their actions; but it seems reasonable to infer that the same condition is placed on the individual’s beliefs. Take the case of *Commonwealth vs. Horsfall*, where a car on a public highway hit an individual, who then died from their injuries. The individual who was killed was stationary, and the driver had seen them from some distance off and sounded their horn. While there was plenty of room to pass, the driver didn’t take it, presumably thinking that the person would move out of the way at the sound of the horn. Even if it might have been rationally permissible to have a high credence that the person would move out of the way, given the driver’s evidence, if that high credence is unusually high or incautious or unreasonable, it seems that its rationality would not exculpate them. Rather, when they are assessed for liability, their action is assessed from the point of view of a person who is cautious in both their beliefs and the actions they perform on the basis of those beliefs.

Now let me explain how we might respect these conservative norms formally. If we wish to extend a credence function in the most conservative way possible, it’s natural to appeal to the Principle of Maximum Entropy (Jaynes, 2003; Paris & Vencovská, 1990, 1997; Williamson, 2010). Typically, that principle applies to an individual whose evidence constrains their credences to some extent, but still permits a range of different credence functions. It is then used to pick out a single credence function from among those: it picks the one that has maximal Shannon entropy.⁹ The idea is this: Shannon entropy measures how unopinionated a probability distribution is. The higher its entropy, the less opinionated it is. Thus, a uniform distribution over a finite partition, which is maximally unopinionated, receives the highest entropy among probability functions over that partition, while a probability function that places all of its mass on a single possible world, and is therefore maximally opinionated, receives the lowest entropy. The idea is that your credence function should respect your evidence; but among credence functions that do this, it should be the least opinionated. In this sense, it should not go beyond the evidence; it should not encode opinions that aren’t demanded by the evidence.¹⁰

In our case, the situation is a little different. It is not only the individual’s evidence that constrains how we might extend their credences to the propositions that lie in \mathcal{F} but not in \mathcal{F}_i . It is also the credences that they assign to the propositions in \mathcal{F}_i . So we might imagine

⁷*Brown vs. Kendall*, 60 Mass. 292 (1850).

⁸*Commonwealth vs. Horsfall*, 213 Mass. 232 (1913).

⁹It is typically used only when it is guaranteed that there will be just one such credence function.

¹⁰It is worth noting that there are other notions of entropy available, and you might prefer to take one of those to measure the unopinionatedness of a probability function instead. I focus here on Shannon entropy, partly because it is the most commonly used, but also because there are persuasive axiomatizations of it, including Shannon’s own (Shannon, 1948; Paris & Vencovská, 1990, 1997).

that each individual has their own body of evidence \mathbf{E}_i , and we might model this as the set of credence functions on \mathcal{F} that respect that evidence. Thus, for instance, if among individual i 's body of evidence is the fact that the coin in their pocket is fair, then each credence function in \mathbf{E}_i should assign credence 50% to that coin landing heads if tossed; and so on. Now, just as we are supposing that all individuals have coherent credence functions, so we might suppose that they all have credence functions that respect their evidence. Thus, for all i , R_i and \mathbf{E}_i overlap. Then we might say: when we extend individual i 's credence function from \mathcal{F}_i to \mathcal{F} , we should ascribe the credence function P_i^{ME} , which is defined on \mathcal{F} as follows:

$$P_i^{\text{ME}} = \arg \sup_{P \in \mathbf{E}_i \cap R_i} H(P)$$

where, recall:

- \mathbf{E}_i is the set of credence functions on \mathcal{F} that respect the evidence that individual i has;
- R_i is the set of coherent credence functions on \mathcal{F} that extend P_i ; and
- $H(P)$ is the Shannon entropy of P .

The motivation is the same as in the standard application of maximal entropy reasoning, where an individual's credences are constrained only by their evidence, and we demand that they pick among those that satisfy the constraints the one that is least opinionated. Similarly here, where both the individual's evidence and their existing credences impose constraints, we ascribe to them the credence function among those that satisfies both constraints that is least opinionated. Thus, we define

$$\Delta_{\text{ME}^*}(P_1, P_2) = \Delta(P_1^{\text{ME}}, P_2^{\text{ME}})$$

where Δ is our favoured pooling function for credence functions defined on the same set of propositions—e.g., linear pooling (Δ_{LP}) or geometric pooling (Δ_{GP}).

Figure 3 illustrates the result of this process in the case we've considered before where:

- each individual has no evidence, so that $\mathbf{E}_1 = \mathbf{E}_2 = \mathcal{P}_{\mathcal{F}}$; and
- the propositions X , Y , and Z form a partition and the individuals' credences are as follows:

	X	Y	Z
P_1	0.8	—	—
P_2	—	0.8	—

Then

	X	Y	Z
P_1^{ME}	0.8	0.1	0.1
P_2^{ME}	0.1	0.8	0.1

Then, if Δ_{LP} is linear pooling, then

	X	Y	Z
$\Delta_{\text{ME}^*}^{\text{LP}}(P_1, P_2)$	0.45	0.45	0.1

It's worth noting that, when we combine this with the illustration from above, we see that taking the credence function that maximises entropy among all linear pools of the possible extensions of P_1 and P_2 is not the same as taking the linear pool of the extensions of P_1 and P_2 that maximise entropy. That is, $\Delta_{\text{ME}}^{\text{LP}}(P_1, P_2) \neq \Delta_{\text{ME}^*}^{\text{LP}}(P_1, P_2)$. And, it seems to me at least, the latter gives the more sensible result.




Figure 3: The barycentric plot of the simplex with $(1, 0, 0)$ at bottom right, $(0, 1, 0)$ at bottom left, and $(0, 0, 1)$ at the top.

9 Conclusion

We've met a lot of different pooling functions that purport to cover those cases in which the individuals in the group in question have different agendas. I have argued that they all fail except Δ_{ME^*} , which I introduced in the previous section. There, I argued that it is the pooling function we ought to use when we wish to determine the corporate credences of a scientific community in order to present them in a textbook, or the corporate credences of a company or institution we are assessing for liability.

Perhaps there are other situations in which it is the pooling function we ought to use, or at least one of the pooling functions we are permitted to use? I think that may well be true when our purpose is not to determine the corporate credences of a group but to determine its aggregate credences. Recall, in Christian List's terminology, the aggregate credences of a group provide a condensed summary of the credences of the individuals that make it up. In this case, there is no suggestion that the group is an agent in its own right. It seems right to say that, when we summarise the credences of a group of individuals, and we need to fill in a particular individual's credence in some proposition in order to perform the summary, we should add as little by way of new opinion as we can. But I don't think this mere appeal to intuition is as convincing as an argument, so I leave this case to future work, when more compelling considerations might be adduced.

10 Appendix: proofs

Proposition 1 *If ∂ is differentiable in its first argument, CAP_∂ violates (EI).*

Proof. Suppose $W = \{w_1, w_2\}$ and $\mathcal{F}_1 = \{w_1, w_2\}$ and $\mathcal{F}_2 = \{w_2\}$. Then (EI) says that, for

any P_1 defined on \mathcal{F}_1 and P_2 defined on \mathcal{F}_2 ,

$$\mathfrak{d}(x, P_1(w_1)) + \mathfrak{d}(1-x, P_1(w_2)) + \mathfrak{d}(1-x, P_2(w_2))$$

is minimized as a function of x at $x = p$ iff

$$\mathfrak{d}(x, P_1(w_1)) + \mathfrak{d}(1-x, P_1(w_2)) + \mathfrak{d}(x, P_2(w_1)) + \mathfrak{d}(1-x, 1-P_2(w_1))$$

is minimized as a function of x at $x = p$. Now, differentiate each with respect to x and evaluate at p , where it will take value 0:

$$\left[\frac{d}{dx} [\mathfrak{d}(x, P_1(w_1)) + \mathfrak{d}(1-x, P_1(w_2)) + \mathfrak{d}(1-x, P_2(w_2))] \right] (p) = \\ \mathfrak{d}'(p, P_1(w_1)) - \mathfrak{d}'(1-p, P_1(w_2)) - \mathfrak{d}'(1-p, P_2(w_2)) = 0$$

And

$$\left[\frac{d}{dx} [\mathfrak{d}(x, P_1(w_1)) + \mathfrak{d}(1-x, P_1(w_2)) + \mathfrak{d}(x, P_2(w_1)) + \mathfrak{d}(1-x, P_2(w_2))] \right] (p) = \\ \mathfrak{d}'(p, P_1(w_1)) - \mathfrak{d}'(1-p, P_1(w_2)) + \mathfrak{d}'(p, P_2(w_1)) - \mathfrak{d}'(1-p, P_2(w_1)) = 0$$

But subtracting the first from the second, we get:

$$\mathfrak{d}'(p, P_2(w_1)) = 0$$

But, since \mathfrak{d} is a divergence, $\mathfrak{d}(x, P_2(w_1))$ is minimized, as a function of x , uniquely at $x = P_2(w_1)$. So $p = P_2(w_1)$. But by similar reasoning, we can also establish:

$$\mathfrak{d}'(p, P_1(w_1)) = 0$$

So it is minimized at $p = P_1(w_1)$. But if $P_1(w_1) \neq P_2(w_1)$, then this gives a contradiction. \square

Theorem 4 is a corollary of this:

Theorem 5. Suppose $f_1, \dots, f_m, g_1, \dots, g_m$ are linear functions in n variables. That is, for each $1 \leq j \leq m$, there are $\alpha_{j1}, \dots, \alpha_{jn}, \beta_{j1}, \dots, \beta_{jn}$ such that

$$f_j(x_1, \dots, x_n) = \alpha_{j1}x_1 + \dots + \alpha_{jn}x_n + k_j$$

and

$$g_j(x_1, \dots, x_n) = \beta_{j1}x_1 + \dots + \beta_{jn}x_n + l_j$$

And suppose $\mathbf{x} = x_1, \dots, x_n$ minimizes

$$\sum_j (f_j(\mathbf{x}) - f_j(\mathbf{p}))^2 + \sum_j (f_j(\mathbf{x}) - f_j(\mathbf{q}))^2 + \sum_j (g_j(\mathbf{x}) - g_j(\mathbf{p}))^2 \quad (1)$$

and $\mathbf{y} = y_1, \dots, y_n$ minimizes

$$\sum_j (f_j(\mathbf{y}) - f_j(\mathbf{p}))^2 + \sum_j (f_j(\mathbf{y}) - f_j(\mathbf{q}))^2 + \sum_j (g_j(\mathbf{y}) - g_j(\mathbf{q}))^2 \quad (2)$$

Then, for $1 \leq i \leq n$, let $z_i = \frac{x_i + y_i}{2}$. Then $\mathbf{z} = z_1, \dots, z_n$ minimizes

$$\sum_j (f_j(\mathbf{x}) - f_j(\mathbf{p}))^2 + \sum_j (f_j(\mathbf{x}) - f_j(\mathbf{q}))^2 + \sum_j (g_j(\mathbf{x}) - g_j(\mathbf{p}))^2 + \sum_j (g_j(\mathbf{x}) - g_j(\mathbf{p}))^2 \quad (3)$$

Proof. Suppose \mathbf{x} minimizes (1). Then, for all $1 \leq k \leq n$,

$$\sum_j 2\alpha_{jk} \left(\sum_i \alpha_{ji}(x_i - p_i) \right) + \sum_j 2\alpha_{jk} \left(\sum_i \alpha_{ji}(x_i - q_i) \right) + \sum_j 2\beta_{jk} \left(\sum_i \beta_{ji}(x_i - p_i) \right) = 0$$

And suppose \mathbf{y} minimizes (2). Then, for all $1 \leq k \leq n$,

$$\sum_j 2\alpha_{jk} \left(\sum_i \alpha_{ji}(y_i - p_i) \right) + \sum_j 2\alpha_{jk} \left(\sum_i \alpha_{ji}(y_i - q_i) \right) + \sum_j 2\beta_{jk} \left(\sum_i \beta_{ji}(y_i - q_i) \right) = 0$$

Now, it's easy to check that, if we let $z_i = \frac{p_i + q_i}{2}$, then \mathbf{z} minimizes (3). So now we need only show that $\frac{x_i + y_i}{2} = \frac{p_i + q_i}{2}$. That is, $x_i = p_i + q_i - y_i$. So, let $x_i^* = p_i + q_i - y_i$. Then

$$\begin{aligned} & \sum_j 2\alpha_{jk} \left(\sum_i \alpha_{ji}(x_i^* - p_i) \right) + \sum_j 2\alpha_{jk} \left(\sum_i \alpha_{ji}(x_i^* - q_i) \right) + \\ & \quad \sum_j 2\beta_{jk} \left(\sum_i \beta_{ji}(x_i^* - p_i) \right) = \\ & \sum_j 2\alpha_{jk} \left(\sum_i \alpha_{ji}((p_i + q_i - y_i) - p_i) \right) + \sum_j 2\alpha_{jk} \left(\sum_i \alpha_{ji}((p_i + q_i - y_i) - q_i) \right) + \\ & \quad \sum_j 2\beta_{jk} \left(\sum_i \beta_{ji}((p_i + q_i - y_i) - p_i) \right) = \\ & \sum_j 2\alpha_{jk} \left(\sum_i \alpha_{ji}(q_i - y_i) \right) + \sum_j 2\alpha_{jk} \left(\sum_i \alpha_{ji}(p_i - y_i) \right) + \\ & \quad \sum_j 2\beta_{jk} \left(\sum_i \beta_{ji}(q_i - y_i) \right) = 0 \end{aligned}$$

Since \mathbf{y} minimizes (2). So \mathbf{x}^* minimizes (1). So $\mathbf{x}^* = \mathbf{x}$, and $x_i = x_i^* = p_i + q_i - y_i$, as required. \square

Conflict of interest There are no conflicts of interest to report.

References

Adamčík, M., & Wilmers, G. (2014). Probabilistic Merging Operators. *Logique et Analyse*, 57(228), 563–590.

- Bamber, J., & Aspinall, W. (2013). An expert judgement assessment of future sea level rise from the ice sheets. *Nature Climate Change*, 3, 424–7.
- Bamber, J., Oppenheimer, M., Kopp, R., Aspinall, W., & Cooke, R. (2019). Ice sheet contributions to future sea-level rise from structured expert judgment. *Proceedings of the National Academy of Science of the U.S.A.*, 166(11195-11200).
- Bradley, R. (2017). *Decision Theory with a Human Face*. Cambridge, UK: Cambridge University Press.
- Bradley, S. (2019). Aggregating Belief Models. *Proceedings of Machine Learning Research*, 103, 38–48.
- Brier, G. W. (1950). Verification of Forecasts Expressed in Terms of Probability. *Monthly Weather Review*, 78, 1–3.
- Bright, L. K., Dang, H., & Heesen, R. (2017). A Role for Judgment Aggregation in Coauthoring Scientific Papers. *Erkenntnis*, 83, 231–252.
- Dang, H. (2019). Do Collaborators in Science Need to Agree? *Philosophy of Science (Proceedings)*, 86(5), 1029–40.
- 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.
- Dietrich, F., & List, C. (2017). Probabilistic Opinion Pooling Generalized. Part One: General Agendas. *Social Choice and Welfare*, 48, 747–86.
- French, S. (1987). Conflict of belief: when advisers disagree. In P. G. Bennett (Ed.) *Analyzing Conflict and its Resolution: Some Mathematical Contributions*, (pp. 93–111). Oxford: Oxford University Press.
- French, S. (2011). Aggregating Expert Judgement. *Revista de la Real Academia de Ciencias Exactas, Fisicas y Naturales. Serie A. Matematicas*, 105, 181–206.
- Galton, F. (1907). Vox Populi. *Nature*, 75, 450–51.
- Genest, C., & Zidek, J. V. (1986). Combining Probability Distributions: A Critique and an Annotated Bibliography. *Statistical Science*, 1(1), 114–135.
- Jaynes, E. T. (2003). *Probability Theory: The Logic of Science*. Cambridge, UK: Cambridge University Press.
- Karni, E., & Vierø, M.-L. (2013). “Reverse Bayesianism”: A Choice-Based Theory of Growing Awareness. *American Economic Review*, 103(7), 2790–2810.
- Konieczny, S., & Pino Pérez, R. (1998). On the logic of merging. In *Proceedings of KR’98*, (pp. 488–498).
- Konieczny, S., & Pino Pérez, R. (1999). Merging with integrity constraints. In *Fifth European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (EC-SQARU’99)*, (pp. 233–244).

- Lackey, J. (2020). *Epistemology of Groups*. Oxford: Oxford University Press.
- List, C. (2014). Three Kinds of Collective Attitudes. *Erkenntnis*, 79, 1601–1622.
- Mahtani, A. (2021). Awareness growth and dispositional attitudes. *Synthese*, 198(8981-8997).
- Osherson, D., & Vardi, M. Y. (2006). Aggregating disparate estimates of chance. *Games and Economical Behavior*, 56(1), 148–173.
- Page, S. E. (2007). *The difference: How the power of diversity creates better groups, firms, schools, and societies*. Princeton, NJ: Princeton University Press.
- Paris, J. B., & Vencovská, A. (1990). A note on the inevitability of maximum entropy. *International Journal of Approximate Reasoning*, 4, 181–223.
- Paris, J. B., & Vencovská, A. (1997). In Defense of the Maximum Entropy Inference Process. *International Journal of Approximate Reasoning*, 17, 77–103.
- Pettigrew, R. (2016). *Accuracy and the Laws of Credence*. Oxford: Oxford University Press.
- Pettigrew, R. (2019). Aggregating incoherent agents who disagree. *Synthese*, 196, 2737–76.
- Predd, J. B., Osherson, D., Kulkarni, S., & Poor, H. V. (2008). Aggregating Probabilistic Forecasts from Incoherent and Abstaining Experts. *Decision Analysis*, 5(4), 177–189.
- Rosenkrantz, R. D. (1981). *Foundations and Applications of Inductive Probability*. Atascadero, CA: Ridgeview Press.
- Shannon, C. E. (1948). A Mathematical Theory of Communication. *Bell System Technical Journal*, 27(3), 379–423.
- Steele, K., & Stefánsson, H. O. (2021). *Beyond Uncertainty*. Cambridge Elements in Decision Theory and Philosophy. Cambridge University Press.
- Tetlock, P. E., & Gardner, D. (2015). *Superforecasting: The Art and Science of Prediction*. Crown Publishers.
- Wenmackers, S., & Romeijn, J.-W. (2016). New theory about old evidence. *Synthese*, 193(4), 1225–1250.
- Williamson, J. (2010). *In Defence of Objective Bayesianism*. Oxford: Oxford University Press.
- Wilmers, G. (2015). A Foundational Approach to Generalising the Maximum Entropy Inference Process to the Multi-Agent Context. *Entropy*, 17, 594–645.