

Groupthink

Jeffrey Sanford Russell, John Hawthorne, Lara Buchak

June 9, 2014

Abstract

How should a group with different opinions (but the same values) make decisions? In a Bayesian setting, the natural question is how to *aggregate credences*: how to use a single credence function to naturally represent a collection of different credence functions. An extension of the standard Dutch-book arguments that apply to individual decision-makers recommends that group credences should be updated by conditionalization. This imposes a constraint on what aggregation rules can be like. Taking conditionalization as a basic constraint, we gather lessons from the established work on credence aggregation, and extend this work with two new impossibility results. We then explore contrasting features of two kinds of rules that satisfy the constraints we articulate: one kind uses fixed prior credences, and the other uses geometric averaging, as opposed to arithmetic averaging. We also prove a new characterisation result for geometric averaging. Finally we consider applications to neighboring philosophical issues, including the epistemology of disagreement.

1 A Problem

The board of Acme Corp is deliberating over whether to invest in an anvil factory. It may succeed, it may fail. If the investment succeeds, the company stands to make 10 thousand dollars profit; if it fails, Acme will lose 11 thousand. But there is disagreement among them over the future of the anvil market. They are evenly divided into two blocs of opinion. One bloc thinks the factory has probability $\frac{2}{3}$ of success; the other puts its chances at only $\frac{1}{3}$. They all agree that the only important thing is to maximize the expected amount of money the company makes. But they have different views on the probabilities of the different outcomes, and none of them has any more say than the others. They realize they have no protocol for resolving whether to take the gamble.

They also know they will face many similar gambles in the future, so they want to settle on a general rule for making this kind of decision. Each of them will state their own credences (with perfect accuracy—they are remarkably good at introspection)

and the rule will somehow aggregate those credences into a single “group credence function” to represent the board collectively, which will dictate their betting behavior.

They start with a simple proposal: since no member has any more say than any other, it seems equitable to take the group credence in each proposition to be the average of their individual credences in that proposition. This will always produce a probabilistically coherent credence function (since each of the individuals has perfectly coherent credences—they are also remarkably internally consistent). Since their average credence that the factory will succeed is $\frac{1}{2}$, the expected value of the gamble is a five hundred dollar loss; so they pass up the opportunity.

Acme Corp also has the opportunity to invest in a balloon factory. The fate of that investment, should it be made, will be decided a year after the fate of the anvils. In this case, they are divided in the same blocs of opinion, but those who are optimistic about anvils are pessimistic about balloons—and vice versa. The pro-balloon bloc thinks this factory has a $\frac{2}{3}$ chance of success, and the pessimists assign it $\frac{1}{3}$. The costs and rewards are the same, so Acme Corp declines this opportunity as well.

Then an enterprising stockbroker approaches them with an offer: a bet on the future of both factories. Acme initially pays 20 thousand dollars. If exactly one of the factories succeed, then they are paid back 37 thousand dollars; otherwise they lose their money.

	B	$\neg B$
A	-20	+17
$\neg A$	+17	-20

Table 1: Net pay-outs for the first gamble.

Everyone on the board agrees that the prospects for the two factories are independent: each person’s credence that anvils boom is the same whether or not balloons bust, and vice versa. So each bloc thinks that this gamble has a $\frac{5}{9}$ chance of paying off. (For the pro-anvil camp, the probability of anvils succeeding and balloons failing is $\frac{2}{3} \cdot \frac{2}{3} = \frac{4}{9}$, and the probability of anvils failing and balloons succeeding is $\frac{1}{3} \cdot \frac{1}{3} = \frac{1}{9}$. For the pro-balloon camp these probabilities are reversed. For both camps the two probabilities sum to $\frac{5}{9}$.) Everyone agrees on this, so the average of their credences is also $\frac{5}{9}$. Since 37 times $\frac{5}{9}$ is more than 20, the expected net return for this gamble is positive (about 560 dollars). It looks like a good move, to each individual and also to the group collectively. So they take the gamble.

A year goes by; anvils do badly. The fate of their gamble now hangs on balloons. The same stockbroker approaches Acme Corp with another proposal, to hedge their potential losses. If they pay 18 thousand dollars now, they will be repaid 37 thousand dollars if the balloon factory fails. Otherwise, they lose their money.

	B	$\neg B$
$\neg A$	-18	+19

Table 2: Net pay-outs for the gamble offered if anvils fail.

The bloc who thought balloons would succeed still have the same opinion: they still think the balloon factory has a $\frac{2}{3}$ chance of success. So this looks like a bad investment to them. But the anti-balloon bloc sees this as a great opportunity, since they still think the balloon factory has only a $\frac{1}{3}$ chance of success. What about the group? The average of their individual credences in the success of balloons is $\frac{1}{2}$. So the group's expected net gain is 500 dollars. Their rules for resolving the disagreement, then, commit Acme Corp to accepting the gamble.

Acme Corp paid a total of 38 thousand dollars for the two gambles. But whether balloons succeed or fail, they will only get 37 thousand dollars back. The stockbroker walks away with a thousand dollars either way.

Furthermore, she had a back-up plan in case anvils succeeded. In that case, their first gamble would only have paid off if the balloons had failed. So she would have offered Acme Corp a chance to hedge their losses by placing another bet, for the same price of 18 thousand dollars, which would pay back 37 thousand dollars if balloons do well. In this case, the pro-balloon bloc would be in favor, the anti-balloon bloc against—and since the average credence that the bet would pay off would again be $\frac{1}{2}$, their rule for resolving their differences would again commit Acme Corp to taking the bet. But in that case, too, the stockbroker makes out like a bandit.

	B	$\neg B$
A	+19	-18

Table 3: Net pay-outs for the gamble offered if anvils succeed.

In fact, even if the stockbroker told the board exactly what she would do in advance, their policy would still commit them to taking the bets and losing money. No individual member would have been bilked this way, but collectively Acme Corp has been diachronically dutch-booked.

The averaging rule got Acme into trouble because of a well-known fact: averaging credences doesn't commute with conditionalization (see for instance Loewer and Laddaga 1985). That is to say, if every individual updates on new information by conditionalizing, the resulting average credence won't be the same as what they would get from first averaging their unconditional credences in each proposition, and then conditionalizing on the new evidence. (We're taking conditionalization to be defined in terms of unconditional credences in the standard way.) This means that even if each individual would take bets like an ideally rational Bayesian agent, the

group won't. And this straightforwardly generalises: any group that fails to conditionalise on new evidence they might receive can fall prey to standard diachronic Dutch books, and so any way of coming up with group credences that does not commute with conditionalisation—like averaging—risks getting the group into trouble.

2 Some Constraints

After this debacle, Acme's board convenes to rewrite their policy. Now that they know that appealing simple policies can get them into trouble, they try a more principled approach, and begin by proposing some general constraints on what a satisfactory aggregation rule would be like.¹ They want a general rule for aggregating their individual credences, no matter what those credences happen to be. This suggests that the rule should be representable as a certain kind of function. We are holding fixed who is in the group and what propositions they are deliberating over; let n be the number of board members and let \mathbf{C} be the set of probability measures on the given algebra of propositions.

Functionality: There is a function $\text{ag} : \mathbf{C}^n \rightarrow \mathbf{C}$ that takes any sequence of probability functions—“the individual credences”—to a probability function—“the group credences”.

(Functionality is a background assumption we will generally hold fixed, and usually we won't mention it in what follows. Note that Functionality by itself is very weak in some respects. For instance, the function is allowed to “build in” an arbitrary amount of detail about the psychology of the individuals, or the intrinsic plausibility of particular propositions. On the other hand it does imply a “Universal Domain” condition: ag is defined for *any* sequence of probability functions. We will consider some natural ways of relaxing Universal Domain later on.)

¹There is a rich mathematical literature on credence aggregation. Genest and Zidek (1986) provide a useful survey of the classic work on this topic. Fitelson and Jehle (2009) present more recent philosophical discussion of some of these results, in the context of the epistemology of disagreement.

This line of inquiry is inspired by parallel results in social choice theory—beginning from Arrow's theorem (1970), which gives an impossibility result for combining preference orderings. This family of results typically involves constraints similar to those we'll discuss, such as Irrelevant Alternatives, Non-Dictatorship, Anonymity, Neutrality, and Unanimity. Arrow's work has also inspired influential work on aggregating “on-off” judgments (for instance, List and Pettit 2002). There is also important work on the more general case of simultaneously aggregating credences and preferences (such as Mongin 1995; Gilboa, Samet, and Schmeidler 2004)—which is not the case we are considering.

The key difference between our work and these earlier results is the prominence we give to Conditionalization, which has no natural analogue in aggregating either preferences or full beliefs, and which (perhaps surprisingly) also has not received much attention in the credence aggregation literature. Two conditions that have received significant attention instead are (Conditional) Independence Preservation and the External Bayesian Condition, which we discuss below. But as the rules we will consider make clear, neither of these conditions are implied by Conditionalization, and so our results go beyond those which appeal to either of them.

(We'll also generally make the simplifying assumption that \mathbf{C} 's algebra of propositions is generated by a finite set of possible worlds, except when otherwise noted. This is not as implausible an assumption as it might seem, since the subject matters for many interesting decision problems like those that face Acme Corp plausibly only distinguish finitely many possibilities. And in any case, it's often useful to start by analyzing the simplest case of a complex question. We'll also assume there are at least three worlds and at least two individuals, to avoid triviality.)

Acme Corp doesn't want to be taken in by any more diachronic Dutch books. So (unlike the averaging rule) ag should commute with conditionalization:

Conditionalization: For any $C = \langle C_1, \dots, C_n \rangle$ in \mathbf{C}^n , if $C' = \langle C_1 | A, \dots, C_n | A \rangle$ is the result of conditionalizing each individual credence function on a proposition A , then $\text{ag } C'$ is the result of conditionalizing $\text{ag } C$ on A .

$$\begin{array}{ccc}
 C & \xrightarrow{\text{conditionalize each } C_i \text{ on } A} & C' \\
 \downarrow & & \downarrow \\
 \text{ag } C & \xrightarrow{\text{conditionalize on } A} & \text{ag } C'
 \end{array}$$

Figure 1: Conditionalization

($C_i | A$ is defined in the usual way, as the credence function that assigns to each proposition B the probability $C_i(B | A) = C_i(A \wedge B) / C_i(A)$. There is a complication: if someone has credence zero in A , then conditionalizing on A is undefined. What the Conditionalization rule should be taken to mean is: if conditionalizing each individual credence on A is *defined*, then so is conditionalizing the group credence function, and it agrees. Note that this has the consequence that if every individual gives a proposition A positive credence, then so does the group. An alternative way to go would be to do everything in terms of primitive *conditional* probabilities, so conditionalizing on a zero-probability proposition might still make sense. But we won't take up the challenge of adapting the results presented here to that setting.)

The board consists of equal partners: no one should have more say than anyone else. It shouldn't make any difference to the rule who has which credences, so if we switch which individual has which credences, the group credence should remain the same:

Anonymity: If C' is any permutation of a sequence C in \mathbf{C}^n , then $\text{ag } C = \text{ag } C'$.

Anonymity is a condition on the "content" of the rule, not how its content was fixed. For instance, if Acme Corp decides to pick a dictator by lot, then though their way of choosing the rule might be intuitively fair, the rule would still not count as Anonymous in this setting. Similarly, the rule might be intuitively unfair while still counting as Anonymous, if one member usurps control of the board and imposes his




Figure 2: Anonymity

own favorite Anonymous aggregation rule—such as one that builds in very opinionated prior credences. (We’ll discuss rules along these lines in Section 3.) Even if these happen to be the usurper’s own credences, this will still count as Anonymous so long as the rule says to use them even in cases where the usurper has different opinions.

Another initially appealing thought is that there should be some simple rule for combining credences in a *single* proposition, which we can then apply to different propositions one by one.²

Systematicity: For any C and C' in \mathbf{C}^n and any propositions A and B , if $C_i(A) = C'_i(B)$ for each i , then $\text{ag } C(A) = \text{ag } C'(B)$.

If this held, then the credence in a proposition wouldn’t depend on anything specific about the proposition besides the numerical value of each individual credence in it. A rule like this would be “topic neutral” in a certain sense. The group credence in a proposition also wouldn’t depend on the credences assigned to other propositions. So the rule would be “local” in a certain sense.

Unfortunately, it has been shown that the *only* kind of aggregation rule that has this property is a weighted average of the individual credences.³ But—generalizing the observation of the previous section—the only way a weighted average rule can obey Conditionalization is if all but one of the weights are zero.⁴ This means that such a rule violates not only Anonymity, but even this weaker condition:

²This principle goes by a variety of names in the literature, including “the strong setwise function property”, “strong label neutrality”, and “the context-free assumption”.

³This was shown independently by McConway (1981) and Wagner (1982). See Genest and Zidek (1986, 117).

⁴Suppose $\text{ag } C$ is the weighted average $\sum_i a_i \cdot C_i$ (with weights a_1, \dots, a_n). Conditionalization tells us

$$\begin{aligned}
\sum_i a_i \cdot C_i(B|A) \cdot C_i(A) &= \sum_i a_i \cdot C_i(A \wedge B) = \text{ag } C(A \wedge B) = \text{ag } C(B|A) \cdot \text{ag } C(A) \\
&= \left(\sum_i a_i \cdot C_i(B|A) \right) \cdot \left(\sum_i a_i \cdot C_i(A) \right)
\end{aligned}$$

In other words, the weighted average of products is the product of weighted averages. This only holds when every weight but one is zero. We can see this by supposing $a_j \neq 0$ and looking at how the group opinion changes when we adjust the the j th conditional credence in B given A , holding everything else fixed. Consider any $C, C' \in \mathbf{C}^n$ where $C_i(A) = C'_i(A) \neq 0$ for all i , $C_j(B|A) = C'_j(B|A)$ for all $i \neq j$, and $C_j(B|A) \neq C'_j(B|A)$. Applying the equation above to both C and C' and taking the difference yields

$$a_j \cdot \left(C_j(B|A) - C'_j(B|A) \right) \cdot C_j(A) = a_j \cdot \left(C_j(B|A) - C'_j(B|A) \right) \cdot \left(\sum_i a_i \cdot C_i(A) \right)$$

Non-Dictatorship: There is no i such that for every C in \mathbf{C}^n , $\text{ag } C = C_i$.

In short: no Non-Dictatorial rule satisfies Systematicity and Conditionalization.

Systematicity combines two ideas: “topic neutrality” and “locality”. We’ve seen that we can’t have both, if our rule is to obey Anonymity and Conditionalization. What if we try to respect just one of the two ideas? Let’s begin by looking at locality: the idea that the group credence in a proposition shouldn’t depend on the individual credences in *other* propositions. This amounts to saying that the group credence in a proposition is a function of the individual credences in that proposition, where the function might vary from proposition to proposition.⁵

Irrelevant Alternatives: For any cases C and C' in \mathbf{C}^n , if $C_i(A) = C'_i(A)$ for each i , then $\text{ag } C(A) = \text{ag } C'(A)$.

Can Irrelevant Alternatives fit with Conditionalization? One reason you might think it doesn’t comes from another important result (Lehrer and Wagner 1983): Irrelevant Alternatives is incompatible with the principle that the group credence **preserves independence** (as long as there are at least five worlds).

Independence Preservation: If A and B are independent according to each individual credence function, then A and B are independent according to the group credence.

But this doesn’t show that Irrelevant Alternatives is incompatible with Conditionalization. There are aggregation rules which obey Conditionalization, but violate Independence Preservation. (We will discuss some examples of such rules in Section 3.) So in general, standard results involving Independence Preservation do not have direct consequences for Conditionalization.⁶ (The same goes for the even stronger principle of *conditional* independence preservation.)

Cancelling non-zero factors, $C_j(A) = \sum_i a_i \cdot C_i(A) = \text{ag } C(A)$. This can only hold generally if j is a dictator, that is, if $a_i = 0$ for all $i \neq j$.

⁵This is also called the “weak setwise function property”, or (confusingly) “Independence”. Here is another equivalent version (McConway 1981; see Genest and Zidek 1986).

Marginalization: For any subalgebra \mathcal{A} of propositions, if sequences of credence functions C and C' agree on \mathcal{A} , then $\text{ag } C$ and $\text{ag } C'$ agree on \mathcal{A} as well.

(The **marginalization** of a credence function is its restriction to a certain subalgebra. So, if you suppose an aggregation rule to be extended to give you a rule that applies to credence functions defined on the subalgebras as well, this principle amounts to saying that aggregation commutes with marginalization.) The thought is that carving up the possibilities more finely, distinguishing more specific subcases, doesn’t make any difference to the group credences in the coarse-grained possibilities.

⁶Independence Preservation is perhaps implausibly strong to begin with. There are cases where two events happen to be independent according to each person’s credences, but intuitively it doesn’t seem important that the group preserve this. Wagner gives this example: if you think a six-sided die is fair, then you should also think that whether an even number is rolled is independent of whether a multiple of three is rolled. Suppose someone else thinks the die is weighted, but in a way that those propositions still happen to come out independent. It’s hard to attach any great importance to keeping this feature of their credences when we combine them. Genest and Wagner (1987) and Wagner (2010b) give further arguments along these lines.

Even so, a different result makes trouble for combining Irrelevant Alternatives with Conditionalization. Note first that Conditionalization implies this:⁷

Zero Preservation: If every individual has zero credence in A , then the group has zero credence in A .

But Wagner also showed that the only kind of rule that obeys Irrelevant Alternatives and Zero Preservation is weighted averaging (see Genest and Zidek 1986, 118). And as before, the only kind of weighted averaging that obeys Conditionalization is a dictatorship.

So it looks like Irrelevant Alternatives is hopeless. To sum up, straightforward applications of well-known results show this:

Fact 1: No rule satisfies Conditionalization, Non-Dictatorship, and Irrelevant Alternatives.

Let's turn now to the other idea involved in Systematicity: topic neutrality. The intuitive thought is that it shouldn't make a difference *which proposition* is assigned which credence. In other words, if we uniformly rearrange individual credences over different worlds, then the group credence should be the result of rearranging the original group credence over the worlds in the same way. Let's put this a bit more precisely. If π is any permutation of the set of worlds and C is a credence function, let πC be the credence function that assigns the same credence to each world πw that C does to w .

Neutrality: For any $C = \langle C_1, \dots, C_n \rangle$ in \mathbf{C}^n and any world-permutation π , the group credence $\text{ag}\langle \pi C_1, \dots, \pi C_n \rangle$ is the same as $\pi(\text{ag } C)$.

$$\begin{array}{ccc}
 C & \xrightarrow{\text{permute worlds}} & \pi C \\
 \downarrow & & \downarrow \\
 \text{ag } C & \xrightarrow{\text{permute worlds}} & \pi(\text{ag } C) = \text{ag}(\pi C)
 \end{array}$$

Figure 3: Neutrality

This property has been less studied in the literature on credence aggregation than the others we have considered so far. One motivation for taking this property as a constraint is the thought that the rule shouldn't "cheat" by consulting outside information beyond what is included in the individual opinions. You ought to be able to read off the group credence from the pattern of individual credences over the different worlds, without knowing what each world represents.

Here is another natural constraint:

⁷This is because $C(A) = 0$ iff $C | \neg A$ is the same as C . (For the right-to-left implication, note that $C(\neg A) = C(\neg A | \neg A) = 1$.) So if $C_i(A) = 0$ for each i , then $\text{ag } C = \text{ag}\langle C_1 | \neg A, \dots, C_n | \neg A \rangle = \text{ag } C | \neg A$ by Conditionalization, and so $\text{ag } C(A) = 0$.

Unanimity: If each individual assigns the same credence to a proposition, then the group also assigns that credence to the proposition.

It would seem strange if every member of the board agreed that a bet on a certain outcome was a bad idea, and yet their rule for resolving disagreements committed them to taking the bet anyhow.⁸

But though these conditions seem natural, another impossibility result follows from them. This result is new, as far as we know.

Fact 2: No rule satisfies Conditionalization, Anonymity, Neutrality, and Unanimity.

The proof turns on a simple symmetry argument. We'll start by considering the simplest non-trivial case, where there are two individuals and three worlds.

Consider a special kind of symmetric case. Suppose that one individual assigns the same credence to w_1 that the other assigns to w_2 , and also vice versa, and that they both assign the same credence to w_3 (Table 4).

	w_1	w_2	w_3
C_1	p_1	p_2	p_3
C_2	p_2	p_1	p_3
$\text{ag}\langle C_1, C_2 \rangle$	q_1	q_2	q_3

Table 4: A symmetric pattern of credences.

In any case with this special structure, the group must assign the same credence to w_1 as it does to w_2 : that is, $q_1 = q_2$. (Let π be the world-permutation that switches w_1 and w_2 . If we apply π to the re-ordered pair $\langle C_2, C_1 \rangle$, the result is equal to the original pair $\langle C_1, C_2 \rangle$. So Anonymity and Neutrality guarantee that applying π to $\text{ag } C$ takes us back to $\text{ag } C$. Since switching w_1 and w_2 leaves $\text{ag } C$ unchanged, the credences in the two worlds must be equal.)

⁸This constraint has the same flavor as Pareto principles—for instance, the one in Arrow's theorem for preference aggregation (see note 1), which says that if every individual ranks X over Y , then the group ranks X over Y as well. (For instance this is how Mongin (1995) motivates Unanimity.)

While we're not sympathetic to Irrelevant Alternatives, it's worth noting as a point of logical geography that Unanimity follows from Irrelevant Alternatives together with a weaker version:

Weak Unanimity: If every individual has the same credences for *every* proposition, then the group also has those credences.

Weak Unanimity is intuitively much weaker: it says nothing at all about what to do in cases of disagreement.

Another point to note is that the following arguments still go through if Unanimity is restricted to apply to cases where the group has "pooled evidence" so each individual assigns the very same propositions credence one, as long as their evidence leaves open at least three worlds.




Figure 4: Symmetry

Furthermore, since in this special case both individuals have the same credence in w_3 , Unanimity guarantees that the group has this credence as well: $q_3 = p_3$. Then, since the probabilities sum to one, for any case with this symmetric structure the constraints uniquely fix the group credences: in fact, in this case they must be the average of C_1 and C_2 . Call this fact **Symmetry**.

But there are cases that are constrained by Symmetry in two different conflicting ways. Here is one pair of credences like that (Table 5). Symmetry leaves us no choice about the group credences here.

	w_1	w_2	w_3
C_1	$1/7$	$4/7$	$2/7$
C_2	$4/7$	$1/7$	$2/7$
$\text{ag}\langle C_1, C_2 \rangle$	$5/14$	$5/14$	$2/7$

Table 5: The credences that make trouble.

	w_1	w_2	w_3
$C_1 A$	0	$2/3$	$1/3$
$C_2 A$	0	$1/3$	$2/3$
$\text{ag}\langle C_1, C_2 \rangle A$	0	$5/9$	$4/9$
$\text{ag}\langle C_1 A, C_2 A \rangle$	0	$1/2$	$1/2$

Table 6: The troublemaking credences after conditionalization.

Now conditionalize on the proposition A that holds at just w_2 and w_3 . This takes the individuals and the group to the credences in Table 6. But in this case, w_2 and w_3 have the special symmetric pattern of individual credences, so Symmetry again requires the group credences to be the *average* of the individual credences—and this gives a different result from conditionalizing the original group credences. So in this case Symmetry contradicts Conditionalization. What this shows is that there are cases C and $C | A$ for which no possible choice of group credences is consistent with all four constraints.

This establishes the result when there are two individuals and three worlds. It is straightforward to generalize the argument. If there are more than three worlds, consider a parallel case where the two individuals have credences proportional to C_1 and C_2 on three worlds, and are completely unanimous on the rest. If there are more than two individuals, consider a case where two individuals have the credences just described, and the rest of the individuals have the average of those two credence functions.

We can also prove a result that does without the assumption of Neutrality, using a different argument.

Fact 3: For a group of two individuals, no rule obeys Conditionalization, Anonymity, and Unanimity.⁹

Conditionalization implies the following fact: the *ratio* of the credences that the group assigns to any pair of worlds is a function of the credence ratios that the individuals assign that pair (when these ratios are all defined). In other words, if there are two cases C and C' , and two worlds v and w such that for each individual i the ratio $C_i(v)/C_i(w)$ is the same as the ratio $C'_i(v)/C'_i(w)$, it follows that $\text{ag } C$ and $\text{ag } C'$ also agree on the credence ratio for those two worlds. (See Lemma 1 in the Appendix.) (It is important that we are restricting our attention to worlds here: Conditionalization does not imply a more general version of this property about the ratio of credences in arbitrary pairs of propositions.¹⁰)

Anonymity guarantees that switching which person assigns which ratio can't make a difference to the group ratio. We will now show that adding Unanimity constrains the ratios still further: if the ratios that two people assign between two worlds v and w are a and b (each of which is at least one), then the group ratio for those worlds must be $\frac{a+b}{2}$. But this constraint is impossible to satisfy in general.

Once again, our strategy is to come up with cases of credences with special symmetries, so they are constrained in more than one way. This time we consider two pairs of credence functions on three worlds (Table 7). Note that C and C' assign the same pair of ratios between w_3 and w_1 (namely $b-1$ and $a-1$). So Conditionalization guarantees that $\text{ag } C$ and $\text{ag } C'$ must also agree on that ratio. Furthermore, the ratios C assigns between w_2 and w_1 are a and b , while the ratios C' assigns between those worlds are b and a . That is, the sequence of individual ratios between w_2 and w_1 given by C just switches the order of the individual ratios given by C' . So Conditionalization and Anonymity together guarantee that $\text{ag } C$ and $\text{ag } C'$ must agree on this ratio as well. But since they agree on the ratio between w_3 and w_1 and also the ratio between w_2 and w_1 , $\text{ag } C$ and $\text{ag } C'$ must be exactly the same. Furthermore, Unanimity tells us that $\text{ag } C$ must assign $\frac{1}{a+b}$ to w_1 and also that $\text{ag } C'$ must assign $\frac{1}{2}$ to w_2 . So the

⁹The argument straightforwardly generalizes to an even number of people, by replacing each individual with a unanimous bloc, but it is less obvious how this would go for an odd number.

¹⁰In fact, no Non-Dictatorial rule can satisfy the more general version—since the more general version implies both Conditionalization and Irrelevant Alternatives (as a special case, considering ratios with a tautology).

ratio of w_2 to w_1 that $\text{ag } C$ and $\text{ag } C'$ both give must be the ratio between these two numbers: that is, $\frac{a+b}{2}$.

	w_1	w_2	w_3
C_1	$\frac{1}{a+b}$	$\frac{a}{a+b}$	$\frac{b-1}{a+b}$
C_2	$\frac{1}{a+b}$	$\frac{b}{a+b}$	$\frac{a-1}{a+b}$
C'_1	$\frac{1}{2b}$	$\frac{1}{2}$	$\frac{b-1}{2b}$
C'_2	$\frac{1}{2a}$	$\frac{1}{2}$	$\frac{a-1}{2a}$

Table 7: The new problem cases.

Moreover, because Conditionalization guarantees that the group ratio between two worlds is a function of the individual ratios, in *any* case where individuals have ratios a and b between w_2 and w_1 (for $a, b \geq 1$), the group ratio must be $\frac{a+b}{2}$. Similar reasoning shows that if the ratios between w_3 and w_2 are both at least one, then the group ratio must be the average ratio; and similarly for w_3 and w_1 . But these three constraints on pairs of worlds can come into conflict. The ratio between w_3 and w_1 is the product of the ratio between w_3 and w_2 and the ratio between w_2 and w_1 . But the product of averages is not generally the average of products. For concreteness, consider this particular pair of credence functions:

	w_1	w_2	w_3
C_1	$\frac{1}{5}$	$\frac{1}{5}$	$\frac{3}{5}$
C_2	$\frac{1}{7}$	$\frac{3}{7}$	$\frac{3}{7}$

Table 8: The average of ratios $w_3 : w_1$ is not the product of the average ratios $w_3 : w_2$ and $w_2 : w_1$.

In this case, the average ratio between w_2 and w_1 is 2 (the average of 1 and 3), and so is the average ratio between w_3 and w_2 . If these were the ratios the group assigned to those pairs, then the group credence in w_3 would have to be 4 times the group credence w_1 . But the average of the individual ratios between w_3 and w_1 is only 2.

(Let's comment briefly on what happens if we drop the assumption that the algebra of propositions is given by a finite set of worlds. The proof of Fact 3 relies on applying the aggregation rule to *discrete* cases, where each individual assigns positive probability to at least two worlds. This makes sense for propositions generated by a countable set of worlds. But in the context of a larger infinite algebra of propositions, discrete probability measures like this might reasonably be thought to be a deviant special case. In that context it is natural to restrict Functionality to exclude these dis-

crete cases from the domain of the aggregation function. With such a restriction, our proof of Fact 3 no longer applies. But Mongin (1995) proves a complementary result, in a rather different way, which only applies to this infinitary context. Suppose that the aggregation rule is restricted to credence functions which are each **non-atomic** in the following sense: for any proposition A with probability p , and any $q \leq p$ there is a proposition B such that $A \wedge B$ has probability q . Mongin shows that in this context Unanimity implies weighted averaging. As we have already discussed, weighted average rules are incompatible with Conditionalization and Anonymity. So our result and Mongin's together show that Fact 3 holds for both the discrete case and the non-atomic case.)

3 Some Aggregation Rules

Now that we have a sense of the constraints on aggregation rules, we'll look at some positive proposals for how to aggregate credences. Each rule we will consider obeys Anonymity and Conditionalization. We've seen that this puts surprisingly strong constraints on what these rules can be like.

Thanks to Fact 1, we know that these rules must violate Irrelevant Alternatives: the group credence in a proposition A must be "holistic" to some extent, depending on other features of the individual credences besides just their opinions about A .

Thanks to Fact 3, we also know that these rules must violate Unanimity: sometimes, the group credence must overrule the unanimous opinions of the individuals.

One further thing to note is that all of these rules violate Independence Preservation: sometimes each individual takes two propositions to be independent, but the group does not. (That means that each of these rules fulfills our promissory note of providing examples where Conditionalization is satisfied but Independence Preservation is not.)

One simple way of getting a rule that obeys Anonymity and Conditionalization is to ignore most features of the group's opinions. We can't ignore *everything*, since Conditionalization implies Zero Preservation. But we can get away with ignoring everything else apart from the facts about which propositions get zero credence. So one thing we can do is at the outset pick some particular prior credence function C^* , and then for any sequence of individual credence functions let the group credence be the result of conditionalizing away the worlds that everyone assigns zero credence. In other words, if E is the conjunction of propositions in which every individual is certain, then the group credence is $C^* | E$. Call this the **Fixed Prior** rule.

(If each individual might assign A positive probability while the prior is zero, then conditionalizing on A would be defined for the individuals but not for the group, violating Conditionalization as we originally stated it. We can avoid this problem by ensuring that every contingent proposition has positive prior probability. This is fine as long as the algebra of propositions is given by a countable set of worlds, as we have been assuming. On the other hand, this approach won't work if this assumption is dropped. If there are uncountably many mutually exclusive propositions, *no* countably

additive credence function can give positive probability to each of them. But there is a natural way of relaxing Functionality—in particular, Universal Domain—that repairs this. Rather than delivering group credences for *arbitrary* individual credence functions, this version of the Fixed Prior rule is only defined for individual credence functions which assign zero probability to each zero-prior proposition. Probability functions like this are called **absolutely continuous**, with respect to the prior. As it turns out, a probability function is absolutely continuous if and only if it can be represented by **probability density**: a function f from worlds to real numbers such that the probability for any proposition A is given by integrating f over A with the prior measure $\mu = C^*$.¹¹ See Figure 5. This density function is uniquely determined up to differences with prior probability zero. The *support* of the density function for an individual credence function, the set of worlds on which f is non-zero, represents that individual’s “evidence”, up to μ -zero differences. Then the Fixed Prior rule says the group credences are given by conditionalizing μ on the conjunction of these evidence-propositions.)




Figure 5: Integrating a density function

Where might this fixed prior come from? It doesn’t have to be any particular person’s prior credences—though it could be. (Choosing a single individual’s prior might seem unfair, but remember that general standards of fairness go beyond the requirements of Anonymity per se.) Or one could randomly generate some credence function, or pick an individual’s prior by lot. Another possibility is to use the average of all of the individual priors. (This doesn’t have to be the arithmetic mean—the alternative kind of averaging we will consider shortly would work just as well.)¹² One more option would be to let C^* be some ideal objective prior, if you believe in such a thing. If you want to know how the corporation *should* bet, a natural interpretation of “should” presses in this direction.

If a Fixed Prior rule is to obey the further constraint of Neutrality—which rules out “cheating” by consulting outside opinions—then C^* must be “uninformative”: in fact, it must be a *uniform* function that spreads credence evenly across the finitely many

¹¹The Radon-Nikodym theorem implies that any absolutely continuous credence function can be represented by a density function this way (see e.g. Halmos 1950, sec. 30–31). Because of the integral formula’s similarity to the fundamental theorem of calculus, the density function f is often called the “derivative” of the measure C and is denoted $\frac{dC}{d\mu}$.

¹²Moss’s (2011) rule would work, too, if it’s applied to the priors. But applied to the posteriors it will violate Conditionalization.

worlds. (If you thought of the ideal objective prior as a kind of Carnapian indifference measure, then the ideal objective prior and the uniform prior might naturally coincide. Note that in the infinite case, there is *no* Neutral prior: there is no probability function on an infinite σ -algebra which is symmetric under arbitrary measurable permutations.¹³)

If we think of having credence one in A as having *evidence* that A , then Fixed Prior effectively conditionalizes on the individuals' *common ground* evidence: the evidence that everyone antecedently shares. An alternative is to instead conditionalize on the individuals' *pooled* evidence, so the corporation takes advantage of the evidence any individual has to offer.¹⁴ In other words, in this alternative version the group credence is $C^* | E$ where E is the conjunction of every proposition that *some* individual is certain of.

The pooled-evidence Fixed Prior rule potentially raises a problem that does not arise for the common ground version: there are sequences of credence functions on which every world is assigned zero credence by *some* individual. In this case, their pooled “evidence” would rule out every world, and so it would be impossible to conditionalize on it. Any way of dealing with this is going to involve relaxing the condition of Functionality, so that at *some* sequences of credence functions the function ag is allowed to be undefined. Suppose having evidence that A is factive. Then these sequences of credence functions, where pooled evidence rules out every world, are simply *impossible* for any individuals to have. (If you are thinking like this then you probably won't think of credence one as a “subjective” state. This also makes the thought that the individuals have access to their own credences look like a more extreme idealization than it already did.) Apart from that, you might not expect there to be anything satisfying to say about what to do if this degenerate case of radical disagreement should arise. We continue to assume that an aggregation rule must be defined in all *other* cases, that is, cases where at least one world is assigned positive credence by every individual. (None of our results relied on assumptions about these degenerate cases, so they still hold using this more relaxed version of Functionality.)

It's clear that Fixed Prior obeys Conditionalization. In either variant, if the “evidence proposition” for certain individual credences is E , and each individual conditionalizes on A , then the new evidence proposition is $A \wedge E$ —and the result of conditionalizing on E and then A is the same as the result of conditionalizing on $A \wedge E$.

There are, however, some prima facie desirable properties that the Fixed Prior rule does not respect. First, as we already noted, Irrelevant Alternatives and Unanimity both fail, as consequences of Facts 1 and 3. It is also clear that this rule does not Preserve Independence. If the prior says A and B are dependent, the group will say

¹³This holds given the standard assumption that coherent credences are countably additive. If this assumption is dropped, there *are* neutral credence functions—for instance, one that assigns probability one to each *cofinite* set of worlds, and zero otherwise. (A set is cofinite if it only leaves out finitely many worlds.)

¹⁴Of course, if we restrict attention to cases where the individuals have *already* pooled their evidence, this version agrees with the common ground version.

this in any case where no individual assign any world zero credence—regardless of their own views on the independence of A and B .

The Fixed Prior rule is also *discontinuous*. An arbitrarily small difference in individual credences can make a large difference to the group credence. In particular, consider a series of cases where one individual’s credence in w approaches zero, while everyone else assigns some constant positive credence. In the pooled evidence version, the group credence has the same positive value in each case—but if that individual’s credence reaches zero, the group jumps down to join her. A similar point holds for the common ground evidence version, where instead each of the other individuals assign zero to w in each case. So either version of the rule has cases involving a discontinuous jump around zero. In short, the Fixed Prior rule violates this:

Continuity: ag is a continuous function.

Finally, there is another notable consequence of the fact that the Fixed Prior rule is not sensitive to any differences in non-zero credences. Even if every individual shifts some of their credence towards a world w , the group will ignore the shift. That fact violates another plausible constraint, which generalizes Conditionalization. Suppose C is a credence function, and η is any function from worlds to non-negative numbers (a **likelihood** function). Let the **η -update of C** be the result of multiplying C by η (pointwise) and renormalizing.¹⁵

External Bayesian Condition: For any η , if C' is the sequence of η -updates of a sequence of individual credence functions C , then $\text{ag } C'$ is the η -update of $\text{ag } C$.

This is equivalent to a condition involving standard Jeffrey conditionalization (R. Jeffrey 1983): if each individual Jeffrey-conditionalizes in a certain way, then the group should Jeffrey-conditionalize in the same way—as long as “same way” is correctly understood. The formulation in terms of likelihoods encourages a particular way of thinking about what updating on the *same* evidence across cases amounts to.¹⁶ Think of your evidence as a set of instructions like “Halve your credence ratio between w_1 and w_2 ”, rather than a set of instructions like “Set your credence in w to $1/3$.” On the first conception of “same evidence” but not the second, two individuals might update on the same evidence without arriving at the same particular credence in any contingent proposition.

¹⁵Note that ordinary conditionalization amounts to the special case where η is zero on some set of worlds and uniform elsewhere. As with Conditionalization, this version of the External Bayesian Condition builds in the assumption that the η -update is well-defined for the group when it is for the individuals.

Again, to generalize this idea beyond the discrete case, it makes sense to restrict attention to absolutely continuous credence functions; then the η -update is given by pointwise multiplication of η with the probability density function.

¹⁶Wagner (2010a) proves the equivalence with Jeffrey conditionalization. See Field (1978) and Wagner (2002) for discussions of the “same evidence” issue.

The Fixed Prior rule is not Externally Bayesian. Suppose everyone has non-zero credences in each of a set of worlds, and the group has the credences Old. The individuals shift their credences in those worlds by some non-uniform likelihood to new non-zero credences. Since the Fixed Prior rule only cares about ones and zeros, the result of aggregating the new credences is exactly the same as Old. On the other hand, the result of updating Old by a non-uniform likelihood is of course not the same as Old.

Let's now turn to another rule which may be a bit less natural to philosophers, but which has been discussed in the statistics literature and does rather better with several of the constraints we have just discussed. It is similar to the original averaging rule, but it uses a different kind of averaging than the simple arithmetic mean. The **geometric mean** of n numbers is the n th root of their product. (Equivalently, the logarithm of the geometric mean is the arithmetic mean of logarithms.) The **Geometric Rule** (or **Logarithmic Rule**) says that the unnormalized group credence in a world w is the geometric mean of the individual credences in w . The group credence in a world, then, is the geometric mean of the individual credences divided by the sum of the geometric means for all worlds.¹⁷

The Geometric Rule forces a pooling approach to evidence: the geometric mean of any number with zero is zero, so if any individual assigns a world zero credence the group must as well. This means, as noted above, the rule is not defined in the degenerate case where every world is assigned zero by at least one individual.

This rule obviously satisfies Anonymity, and it also satisfies Conditionalization—in fact, it obeys the stronger External Bayesian Condition. The general version follows from the fact that if you multiply some numbers by a common factor $\eta(w)$ and then take their geometric mean, this gives you the same result as if you take the geometric mean first and then multiply by $\eta(w)$. So it doesn't make a difference whether you η -update and then aggregate, or do it in the opposite order. Since the Geometric Rule obeys Anonymity and Conditionalization, it follows from Fact 3 that it must not satisfy Unanimity. The reason is that, even though taking the geometric mean preserves unanimity, renormalizing the geometric mean may not.¹⁸

Let's look at how this plays out for Acme Corp's original diachronic Dutch book.

¹⁷This rule is attributed to Peter Hammond, who also noted the fact that it obeys the External Bayesian Condition (Genest and Zidek 1986, 119–20). In a blog post (2012) Alexander Pruss makes a closely related suggestion for aggregating credences in a single proposition from individuals with the same evidence (namely, averaging the logarithm of odds), and discusses some of its nice features and an alternative motivation.

¹⁸On the other hand, the Geometric Rule does obey these:

Pointwise Ratio Unanimity: For any pair of worlds w_1 and w_2 , if each individual in C assigns the same credence ratio between w_1 and w_2 , then $\text{ag } C$ also assigns that ratio.

Pointwise Comparative Unanimity: If each individual assigns a higher credence in world w_1 than w_2 , then the group does as well.

As with Unanimity, there are natural analogies between these and the Pareto principle—though there is also a disanalogy, in that these only apply “world by world”. The Geometric Rule does not satisfy the more general versions for arbitrary propositions.

The stockbroker’s second proposal to Acme, after the anvils had failed, was a bet with a net payout of 19 thousand dollars if balloons failed, and a net loss of 18 thousand dollars if balloons succeeded. The board was divided on this, one bloc giving the happy outcome a probability of $\frac{1}{3}$ and the other bloc giving it $\frac{2}{3}$. The renormalized geometric mean of these credences is $\frac{1}{2}$. So the Geometric Rule tells the group to take the bet. (The same reasoning applies to the alternative second bet.) What about the first bet? Recall that this was a bet with a net payout of 17 thousand if exactly one of anvils and balloons succeeded, and a net loss of 20 thousand otherwise. Their credences were as in Table 9. In this case, each individual favored the gamble. But the Geometric Rule overrules their unanimous opinion—this is how the Geometric Rule saves Acme Corp from being Dutch-booked. As we have shown already, any rule that obeys Conditionalization must overrule unanimous opinions somewhere.¹⁹

	$A \wedge B$	$A \wedge \neg B$	$\neg A \wedge B$	$\neg A \wedge \neg B$
Pro-anvil	$\frac{2}{9}$	$\frac{4}{9}$	$\frac{1}{9}$	$\frac{2}{9}$
Pro-balloon	$\frac{2}{9}$	$\frac{1}{9}$	$\frac{4}{9}$	$\frac{2}{9}$
Geometric Rule	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$

Table 9: Applying the Geometric Rule to Acme Corp’s original credences.

Note that the Geometric Rule does not say that the group credence in an arbitrary proposition is given by the geometric mean of individual credences (followed by some renormalization)—it only applies directly to worlds. How credences are distributed over subcases of a proposition can make a difference. (This is a consequence of the fact that this rule violates Irrelevant Alternatives—which follows from Fact 1.) Note, for example, that for Acme Corp’s first bet each individual thought that the probability of exactly one factory succeeding was $\frac{5}{9}$. So if we had simply applied the Geometric Rule directly to the individual credences in the proposition “exactly one factory succeeds” and its negation, ignoring their differences of opinion on subcases, then we would have got a different result—the renormalized geometric mean would have been $\frac{5}{9}$, and the rule would recommend taking the bet. But “coarse-graining” their credences this way would wipe out all of the information which is relevant to the later bets on each one of the factories.

The Geometric Rule also violates Independence Preservation. (A result of Genest and Wagner (1987, 82–3) shows that it in fact Preserves Independence as long as there are no more than four worlds, but not otherwise.)

There are natural generalizations of the Geometric Rule that give up Anonymity or Neutrality, by respectively assigning non-uniform weights to particular individu-

¹⁹Naturally this goes for the Fixed Prior rule, too, but the details of its recommendations will vary depending on what the fixed prior is. Some versions will reject the first bet, and others will reject one of the second bets.

als or to particular worlds. To give different weights to different individuals, we can use a weighted geometric mean—here the weights appear as *powers* in the product of probabilities. The geometric mean of p_1, \dots, p_n with weights a_1, \dots, a_n , is the product $p_1^{a_1} \cdot \dots \cdot p_n^{a_n}$. Call rules of this form **Weighted Geometric Rules**: take a pointwise weighted geometric mean, and renormalize. The unweighted Geometric Rule is the special case of this where each weight is $1/n$.

To give different weights to different worlds—which you can think of as treating different worlds as building in some opinionated prior probabilities—we can introduce an extra credence function as one more factor to this product, with some weight of its own. This is perhaps the best way of generalizing the Fixed Prior rule in a way that is sensitive to more features of the individual credences than just ones and zeros.

We already noted that unlike Fixed Prior, the Geometric Rule is Externally Bayesian. Another advantage it has over Fixed Prior is that it is Continuous. In fact, we can characterize the rule this way.

Fact 4: The only rules which obey Conditionalization, Continuity, and Neutrality are Weighted Geometric Rules.

Since the proof of this fact is a bit more technical than the others, we present it in an appendix.²⁰ If we add Anonymity as a further constraint, this forces the weights to be equal. If we add the Weak Unanimity constraint—which says that in cases where every individual has exactly the same credences about everything, the group has those too—then this forces the weights to add up to one. So if we add both of those constraints each weight must be $1/n$, so we have the standard Geometric Rule. That is to say, the only rule which obeys Conditionalization, Continuity, Neutrality, Anonymity, and Weak Unanimity is the Geometric Rule.

(The Geometric Rule extends straightforwardly to the case of countably many worlds.²¹ But the more general uncountable case is more complicated: in this case arbitrary credences aren't determined by credences in particular worlds, and so pointwise geometric averaging doesn't determine group credences. But there is a natural, and technically standard, extension of the *opinionated* version of the rule, which uses the same ideas as the infinitary generalization of the Fixed Prior rule. The idea is

²⁰This result complements those of Genest (1984) and Genest, McConway, and Schervish (1986). Genest (1984) shows that weighted geometric averaging is the only kind of rule that is Externally Bayesian and also obeys a weakened form of Irrelevant Alternatives. (*Viz*: the group's probability density at a world is determined by the individual's probability densities at that world, up to a constant normalization factor.) Genest, McConway, and Schervish (1986) extend this result to a general characterization of Externally Bayesian operators. The most important difference between our result and these is that we do not rely on the External Bayesian Condition, but only the weaker Conditionalization principle. Also, our Neutrality and Continuity conditions are orthogonal to Genest's Irrelevant-Alternatives-style principle. Finally, the context our result applies to countable probability measures, rather than probability *densities* in a more general setting.

²¹This turns on the fact that if the sums of two infinite sequences converge, then the sum of the sequence of their geometric means also converges. (This is clear, since this sequence is bounded by the pointwise *maximum* of the two sequences, and the sum of the maximums must converge.)

to first fix some “background” measure μ .²² Functionality is restricted, so the domain of the rule only includes individual credences which are absolutely continuous with respect to μ . This guarantees that the individual credences are represented by probability densities. So we can apply the Geometric Rule by taking the pointwise geometric mean of the density functions, and renormalizing, to get a group density. Note that the Neutral version of the Geometric Rule for countably many worlds is the special case of this where the background measure μ is the counting measure, which gives each world equal weight.)

4 Some Connections

We’ve been focusing on the problem of how groups of individuals can collectively have coherent credences. We’ll conclude by pointing to some applications of these ideas to other philosophical topics. One natural connection is to another issue in social epistemology, namely how individuals should update their credences in response to disagreements with their peers. Some philosophers have defended the view that disagreeing peers should give “equal weight” to each person’s credences regardless of whose they are, and adjust their own credences to some value that impartially reflects all of the disagreeing opinions (Elga 2007; Christensen 2007; for critical discussion see Lackey 2008; Kelly 2010). In order to make sense of a view like this, though, it is necessary to have some idea of what those impartially generated credences would be. Both the limiting and positive results about aggregation functions have implications for the shape that this sort of view can take. (See also Fitelson and Jehle 2009; Moss 2011.)

Normally people are not “peers” about everything, but at best some distinguished subject matter—the **peer propositions**. These presumably will not be an arbitrary set: if they are peers about A and peers about B , then they are also peers about $\neg A$, $A \wedge B$, and $A \vee B$. Then if we restrict each peer’s credence function to the peer propositions, we get a credence function on this special subject matter—call these restricted credence functions the **peer credences**. Then the idea of the Equal Weight view is that there should be some way of generating **impartial credences** from the peer credences, which are what each peer rationally ought to adopt. (Note that this set-up assumes, for better or worse, that the impartial credence in a *peer* proposition does not depend on any individual’s credences in a *non-peer* proposition. The set-up also leaves open the question of how peers ought to adjust the rest of their credences. One natural thought is that they should Jeffrey conditionalize on the impartial credences—there is a technical sense in which this is the “minimal” adjustment.)

The most obvious Equal Weight view is one that says the impartial credence is the arithmetic mean of the peer credences. As we discussed in Section 1 this rule violates Conditionalization. This failure is straightforwardly bad for betting board members, but there is some debate about how serious it is here (see for instance Wilson 2010).

²²This need not be a probability measure, but it should at least be σ -finite, meaning that the worlds can be partitioned into countably many pieces of finite measure.

Those who do think Conditionalization is an important constraint should also take to heart the more general point (Fact 1) that no Conditionalizing rule besides dictatorships can also obey Irrelevant Alternatives. There is a general tendency in these discussions to take Irrelevant Alternatives for granted by assuming that it makes sense to consider peer propositions one by one, where the impartial view on any proposition is determined by the peer credences in just that one. One lesson from the impossibility results is that, if you care about Conditionalization, you should be thinking holistically about the whole peer subject matter rather than just single propositions. Discussions also tend to assume that Unanimity is an important constraint: peers who already agree on A shouldn't move away from that credence when they learn of their disagreement on other matters. But as we have seen, if you care about Conditionalization then this principle is also difficult to sustain.

The aggregation rules we have discussed also give guidance on how to avoid some of these technical problems. For example, the Geometric Rule might be a much better candidate than simple averaging for what “splitting the difference” between credence functions ought to amount to. It implements a reasonable notion of impartiality, while still respecting Conditionalization. Similarly, an approach that averages peer priors and then updates on peer evidence does better.

Besides social epistemology, aggregation issues also arise for individual epistemology insofar as a person can be “double-minded” in various ways. There are natural approaches to imprecise or “mushy” credences, higher-order uncertainty about one's own credences, and psychological fragmentation, which involve representing a single person's epistemic state by a family of distinct credence functions (for instance, Levi 1980; R. C. Jeffrey 1983; for critical discussion based on concerns related to ours, see Elga 2010). Despite this multitude of opinions, we want to say something about how a fragmented person can take a unified rational stance on gambles, which would seem to require some way of aggregating the fragments.

There are also potential applications to political theory. As we mentioned in footnote 1, there is a well-known body of “voting theorems” constraining how individual preferences can be aggregated, and these results have been extensively applied to the theory of democratic government. Constraints on credence aggregation have some analogous implications, since of course some political disagreements are naturally represented as conflicts of *beliefs* rather than conflicts of *values*.

A Proof of Fact 4

Let W be a countable set of worlds. In this context, a credence function is given by a function from W to $[0, 1]$ that sums to one (i.e., a probability mass function). Let n be the number of individuals. Call a sequence of n credence functions **admissible** iff there is some world that each function gives positive probability. (This restriction goes with the idea discussed in Section 3 that credence one is factive.) Then the aggregation rule **ag** is a function that takes each admissible sequence to a single credence function. In this setting, **Conditionalization** says that for any sequence $C = \langle C_1, \dots, C_n \rangle$

and set of worlds E , if $C | E = \langle C_1 | E, \dots, C_n | E \rangle$ is defined and admissible, then $\text{ag}(C | E) = \text{ag } C | E$. In what follows C and C' are admissible sequences of credence functions.

For a credence function C_i , let $C_i(v : w)$ stand for the ratio $C_i(v)/C_i(w)$ (if it is defined).

Pointwise Ratios: For any $v, w \in W$, if for each i , $C_i(v : w)$ and $C'_i(v : w)$ are defined and equal, then $\text{ag } C(v : w) = \text{ag } C'(v : w)$.

Lemma 1. Conditionalization is equivalent to Pointwise Ratios.

Proof. Let $E = \{v, w\}$. For each i , if $C_i(v : w) = C'_i(v : w)$ then $C_i | E = C'_i | E$. So by Conditionalization, $\text{ag } C | E = \text{ag } C' | E$. This implies that $\text{ag } C(v : w) = \text{ag } C'(v : w)$ as well. So Conditionalization implies Pointwise Ratios.

Conversely, let E be any proposition such that $C | E$ is defined and admissible, and let v be a world in E that each individual gives positive probability. Then for each $w \in E$, C and $C | E$ both have the same well-defined ratios between w and v . So $\text{ag } C$ and $\text{ag}(C | E)$ also have the same ratios for each $w \in E$. So $\text{ag } C | E$ and $\text{ag}(C | E)$ are proportional, and since each of them adds up to one they are identical. So Pointwise Ratios implies Conditionalization. QED.

Recall that **Neutrality** means that ag commutes with permutations of W , and **Continuity** means that ag is a continuous function (with respect to the product topology of $[0, 1]^{W \times \mathcal{A}}$).

Fact 4: If ag obeys Conditionalization, Continuity, and Neutrality, ag is a Weighted Geometric Rule.

Proof. Pointwise Ratios and Neutrality together imply that there is some function $F : [0, \infty)^n \rightarrow [0, \infty)$ such that for each pair of distinct worlds v and w , if the ratios $C_1(v : w), \dots, C_n(v : w)$ are all defined, then

$$\text{ag } C(v : w) = F(C_1(v : w), \dots, C_n(v : w))$$

(Pointwise Ratios guarantees that for each v and w there is some function $F_{v,w}$ that determines the group ratio for v and w in terms of the individual ratios, when they are defined. Neutrality guarantees that $F_{v,w}$ is the same for each v and w .) Furthermore, if ag is continuous then F is continuous as well.

Ratios have the following property: if $C(u : v)$ and $C(v : w)$ are both defined, then $C(u : w) = C(u : v) \cdot C(v : w)$. This implies that F is *multiplicative* for positive arguments:

$$F(r_1 \cdot s_1, \dots, r_n \cdot s_n) = F(r_1, \dots, r_n) \cdot F(s_1, \dots, s_n)$$

(where each r_i and s_i is positive.) It's helpful to map this onto a logarithmic scale: there is a continuous function $G : \mathbb{R}^n \rightarrow \mathbb{R}$ such that

$$G(\log r_1, \dots, \log r_n) = \log F(r_1, \dots, r_n)$$

(for positive r_i). It follows from F 's multiplicative property that G is *additive*:

$$G(x_1 + y_1, \dots, x_n + y_n) = G(x_1, \dots, x_n) + G(y_1, \dots, y_n)$$

But any continuous additive function from \mathbb{R}^n to \mathbb{R} is *linear*. (This fact was noted by Cauchy in 1821.) So G is linear, and thus there are weights a_1, \dots, a_n , such that

$$G(x_1, \dots, x_n) = a_1 \cdot x_1 + \dots + a_n \cdot x_n$$

Undoing the transformation to the logarithmic scale, then,

$$F(r) = r_1^{a_1} \cdot \dots \cdot r_n^{a_n}$$

This fixes the value of F for positive ratios to be a weighted geometric mean. When $r_i = 0$, continuity forces F 's value be the limit value as r_i approaches zero. Accordingly, if $a_i > 0$, then $F(r_1, \dots, r_n)$ must be zero when $r_i = 0$, which is consistent with geometric averaging. For $a_i = 0$, if we consider $0^0 = 1$ then again the geometric average extends F continuously to $r_i = 0$. For $a_i < 0$ there is no finite limit at zero, so that case is impossible for continuous F . So F —the rule for group *ratios*—is a weighted geometric mean with non-negative weights.

This means that ag is a Weighted Geometric Rule. Let v be a world with positive individual credences p_1, \dots, p_n , and say the group credence in that world is p . Then for any other world w with individual credences q_1, \dots, q_n the ratios q_i/p_i are defined, and the group ratio is a weighted geometric mean of those ratios. So the group credence in w is

$$p \cdot \left(\frac{q_1}{p_1}\right)^{a_1} \cdot \dots \cdot \left(\frac{q_n}{p_n}\right)^{a_n}$$

which, redistributing parentheses, is just the weighted geometric mean $q_1^{a_1} \cdot \dots \cdot q_n^{a_n}$ times a constant normalization factor. QED.

References

- Arrow, Kenneth J. 1970. *Social Choice and Individual Values, Second Edition*. Yale University Press.
- Christensen, David. 2007. "Epistemology of Disagreement: the Good News." *Philosophical Review* 116 (2): 187–217.
- Elga, Adam. 2007. "Reflection and Disagreement." *Noûs* 41 (3): 478–502.
- . 2010. "Subjective Probabilities Should Be Sharp." *Philosophers' Imprint* 10 (05).
- Field, Hartry. 1978. "A Note on Jeffrey Conditionalization." *Philosophy of Science* 45 (3): 361–367.
- Fitelson, Branden, and David Jehle. 2009. "What Is the 'Equal Weight View'?" *Episteme* 6 (3): 280–293.

- Genest, Christian. 1984. "A Characterization Theorem for Externally Bayesian Groups." *The Annals of Statistics* 12 (3) (September): 1100–1105. <http://www.jstor.org/stable/2240984>.
- Genest, Christian, and Carl G. Wagner. 1987. "Further Evidence Against Independence Preservation in Expert Judgement Synthesis." *Aequationes Mathematicae* 32 (1): 74–86. <http://www.springerlink.com/index/01661U7611117843.pdf>.
- Genest, Christian, and James V. Zidek. 1986. "Combining Probability Distributions: a Critique and an Annotated Bibliography." *Statistical Science* 1 (1) (February): 114–135. <http://www.jstor.org/stable/2245510>.
- Genest, Christian, Kevin J. McConway, and Mark J. Schervish. 1986. "Characterization of Externally Bayesian Pooling Operators." *The Annals of Statistics* 14 (2) (June): 487–501. <http://www.jstor.org/stable/2241231>.
- Gilboa, Itzhak, Dov Samet, and David Schmeidler. 2004. "Utilitarian Aggregation of Beliefs and Tastes." *Journal of Political Economy* 112 (4) (August): 932–938. doi:10.1086/421173. <http://www.jstor.org/stable/10.1086/421173>.
- Halmos, Paul R. 1950. *Measure Theory*. Princeton: D. Van Nostrand.
- Jeffrey, Richard. 1983. *The Logic of Decision*. University of Chicago Press.
- Jeffrey, Richard C. 1983. "Bayesianism with a Human Face." In *Testing Scientific Theories*. University of Minnesota Press.
- Kelly, Thomas. 2010. "Peer Disagreement and Higher Order Evidence." In *Social Epistemology: Essential Readings*. Oxford University Press.
- Lackey, Jennifer. 2008. "What Should We Do When We Disagree?" In *Oxford Studies in Epistemology*. Oxford University Press.
- Lehrer, Keith, and Carl G. Wagner. 1983. "Probability Amalgamation and the Independence Issue: a Reply to Laddaga." *Synthese* 55 (3): 339–346. doi:10.1007/BF00485827. <http://www.springerlink.com/content/k6vh273608174u98/abstract/>.
- Levi, Isaac. 1980. *The Enterprise of Knowledge: an Essay on Knowledge, Credal Probability, and Chance*. The Mit Press.
- List, Christian, and Philip Pettit. 2002. "Aggregating Sets of Judgments: an Impossibility Result." *Economics and Philosophy* 18 (1): 89–110.
- Loewer, Barry, and Robert Laddaga. 1985. "Destroying the Consensus." *Synthese* 62 (1): 79–95.
- McConway, Kevin J. 1981. "Marginalization and Linear Opinion Pools." *Journal of the American Statistical Association* 76 (374) (June): 410–414. doi:10.2307/2287843. <http://www.jstor.org/stable/2287843>.
- Mongin, Philippe. 1995. "Consistent Bayesian Aggregation." *Journal of Economic Theory* 66 (2) (August): 313–351. doi:10.1006/jeth.1995.1044. <http://www.sciencedirect.com/science/article/pii/S0022053185710447>.
- Moss, Sarah. 2011. "Scoring Rules and Epistemic Compromise." *Mind* 120 (480): 1053–1069.
- Pruss, Alexander. 2012. "Aggregating Data from Agents with the Same Evidence." *Alexander Pruss's Blog*. <http://alexanderpruss.blogspot.co.uk/2012/03/aggregating-data-from-agents-with-same.html>.

- Wagner, Carl G. 1982. "Allocation, Lehrer Models, and the Consensus of Probabilities." *Theory and Decision* 14 (2): 207–220. doi:10.1007/BF00133978. <http://www.springerlink.com/content/u186nm5gxp990221/abstract/>.
- . 2002. "Probability Kinematics and Commutativity." *Philosophy of Science* 69 (2) (June): 266–278. <http://www.jstor.org/stable/10.1086/341053>.
- . 2010a. "Jeffrey Conditioning and External Bayesianity." *Logic Journal of IGPL* 18 (2): 336–345. <http://jigpal.oxfordjournals.org/content/18/2/336.short>.
- . 2010b. "Peer Disagreement and Independence Preservation." *Erkenntnis* 74 (2) (November): 277–288. doi:10.1007/s10670-010-9256-9. <http://www.springerlink.com/index/10.1007/s10670-010-9256-9>.
- Wilson, Alastair. 2010. "Disagreement, Equal Weight and Commutativity." *Philosophical Studies* 149 (3).