

Episteme

<http://journals.cambridge.org/EPI>


Additional services for *Episteme*:

Email alerts: [Click here](#)

Subscriptions: [Click here](#)

Commercial reprints: [Click here](#)

Terms of use : [Click here](#)



A NEW PROSPECT FOR EPISTEMIC AGGREGATION

Daniel Berntson and Yoaav Isaacs

Episteme / Volume 10 / Issue 03 / September 2013, pp 269 - 281

DOI: 10.1017/epi.2013.26, Published online: 07 August 2013

Link to this article: http://journals.cambridge.org/abstract_S1742360013000269

How to cite this article:

Daniel Berntson and Yoaav Isaacs (2013). A NEW PROSPECT FOR EPISTEMIC AGGREGATION. Episteme, 10, pp 269-281 doi:10.1017/epi.2013.26

Request Permissions : [Click here](#)

A NEW PROSPECT FOR EPISTEMIC AGGREGATION

DANIEL BERNTSON AND YOAAV ISAACS*

berntson@princeton.edu and yisaacs@princeton.edu

ABSTRACT

How should the opinion of a group be related to the opinions of the group members? In this article, we will defend a package of four norms – coherence, locality, anonymity and unanimity. Existing results show that there is no tenable procedure for aggregating outright beliefs or for aggregating credences that meet these criteria. In response, we consider the prospects for aggregating credal pairs – pairs of prior probabilities and evidence. We show that there is a method of aggregating credal pairs that possesses all four virtues.

I. A MURDER IN GRAND FENWICK

There has been a murder of such staggering significance that it could only take place in fiction. Leopold, the Archduke of Grand Fenwick, lies dead in a pool of blood, his royal saber buried in his back. To investigate this heinous crime, the greatest detectives from any nation, any time and any genre have been assembled: Sherlock Holmes, Hercule Poirot and Jane Marple. The great detectives agree (as anyone with experience of fictional crimes would) that motive and opportunity are individually necessary and jointly sufficient conditions for murder. The unique person with both motive and opportunity for Archduke Leopold's murder is Archduke Leopold's murderer.

Suspicion naturally falls upon Otto, Leopold's younger brother, the next in line for the throne of Grand Fenwick. Given the threat of a political crisis, Leopold's chief of staff orders the detectives to investigate Otto. Depending on the detectives' report, Otto will either be executed or ascend to the throne. The detectives swiftly conduct their investigation, and are soon called upon by the chief of staff to report their findings. Holmes replies, "Otto had motive but not opportunity. Otto is not the murderer." Poirot replies, 'Otto had opportunity but not motive. Otto is not the murderer.' Marple replies, 'Otto had both motive and opportunity. Otto is the murderer.' The chief of staff is unimpressed. 'Come now' he says, 'I did not ask for each of your opinions. I asked for your joint opinion. I can take only one course of action regarding Otto, and so I can accept only one opinion from the three of you. Tell me: what do you – the three of you – think?'

* As this article represents the equal contributions of both authors, any praise and/or blame for what follows should be distributed equally. We would also like to thank Thomas Kelly, Sarah-Jane Leslie, Branden Fitelson, Eric Hubble, an anonymous referee and the participants at the 2013 Penn-Rutgers-Princeton Social Epistemology Workshop for especially helpful comments and questions.

At this, Holmes, Poirot and Marple were silent. They each knew perfectly well what they thought as individuals, but they had no idea what they thought as a group.

2. EPISTEMIC AGGREGATION

Holmes, Poirot and Marple face the problem of epistemic aggregation – how to combine individual opinions into a single, group opinion.

Consider some of the possibilities for aggregating the beliefs of the three great detectives.¹ One idea is to say that a group believes a proposition just in case all the individuals believe the proposition. Call this aggregation procedure **intersection**. In our case, the detectives are unanimous when it comes to certain logical truths (Otto either had motive or did not have motive) and when it comes to some general features of the case (Otto either had the opportunity or the motive to kill the Archduke). But since they are not unanimous when it comes to, say, whether or not Otto committed the murder, the group will simply have no opinion when it comes to the most important features of the case.²

Here is another suggestion. Say that the group of detectives collectively believe a proposition just in case a majority of the detectives believes that proposition. But this method has a troubling consequence. Two out of three detectives believe that Otto had motive, two out of three believes that he had opportunity, and two out of three believe that he did not have both motive and opportunity. Using this method of epistemic aggregation, it turns out that the group believes that Otto had motive, that Otto had opportunity, and that Otto did not have motive *and* opportunity. Three consistent epistemic states are aggregated into an inconsistent epistemic state.³

Not only is the aggregation of beliefs difficult, we can in fact prove that intersection is the only aggregation procedure that has certain important and desirable features (for finite groups of individuals at least).

The first feature is **coherence**.⁴ If all the individuals in a group have logically consistent beliefs, then the group should have logically consistent beliefs as well.

1 These possibilities are relevant both for descriptive questions in social epistemology (about what groups *believe*) and for normative questions in social epistemology (about what groups *should* believe). Our main interest is in determining what sorts of epistemic aggregation are possible, and specifically whether any procedure for epistemic aggregation can possess certain desirable features.

2 Even if you think that this is the right result in the present case, the requirements of intersection are almost certainly too strong in other cases. Suppose that 1,000 detectives look at the case and that all but one agree that Otto committed the murder: it would be odd to think that the group has no opinion on the matter.

3 There is reason to think that logical consistency is not a good norm for outright beliefs. Preface and lottery cases suggest that it can be rational to hold sets of beliefs that are inconsistent, at least when those sets are sufficiently large (if there is a lottery with a million tickets, it is rational to believe of each ticket that it will lose and to believe that some ticket will win even though these beliefs are logically inconsistent). There nonetheless seems to be something defective about the inconsistency of such small sets of beliefs. Since for our purposes the aggregation of outright belief is primarily illustrative, we can set aside these issues and think of logical consistency as the norm for coherent beliefs. For a good discussion of various candidate coherence norms for outright belief, see Fitelson and Easwaran (forthcoming).

4 See Pettit and List (2004) for more on the difficulties of coherently aggregating outright beliefs (as well as for the inspiration for our three detectives).

A second feature it would be nice for an aggregation procedure to have is **locality**. The intuitive idea behind locality is that what the group believes about a proposition p is fully determined by what the individuals think about p . We can state this constraint a bit more precisely by thinking of individual beliefs as characteristic functions $B^i(\cdot)$ (these functions return 1 if the individual i believes a proposition and 0 if she does not) and aggregation as a relation between sequences of belief states $S = \langle B^1, B^2 \dots \rangle$ and group belief states B^* . Locality then says that there is some function $f: \{0, 1\}^n \rightarrow \{0, 1\}$ such that B^* is an aggregate belief state of $\langle B^1, B^2 \dots \rangle$ just in case $f(B^1(p), B^2(p) \dots) = B^*(p)$ for all propositions p .⁵

The third desirable feature is **anonymity**. When it comes to the three great detectives, anonymity says that if we were to rearrange the opinions of the individuals in the group – if Holmes had Poirot’s beliefs, Poirot had Marple’s beliefs and Marple had Holmes’ beliefs – the aggregate opinion of the group would remain the same. Put more formally, if we have a sequence of belief states S , permuting S and then aggregating the result yields the same output as aggregating S .

The final feature is **unanimity**. Unanimity says that whenever all the members in a group believe that p , the group also believes that p .

We can now prove that intersection is the only aggregation procedure that has all of the above features. Locality tells us that what the aggregate believes about p is determined by only what the individuals believe about p , and anonymity says that we can permute who thinks what about p without changing what the group thinks about p . Locality also tells us that if a certain pattern of belief is sufficient for belief in p , then that pattern is also sufficient for belief in any other proposition q . In particular, if the group believes p when n out of m individuals believe p , then the group also believes q when n out of m individuals believe q . Now suppose that n is less than m . We can construct a lottery paradox-like case in which, for each of n propositions p_n , n out of m individuals believe p_n and all of them believe the negation of the conjunction $p_1 \wedge \dots \wedge p_n$. By unanimity, the group then also believes the negation of the conjunction $p_1 \wedge \dots \wedge p_n$. But then the group has a logically inconsistent belief state, even though all of the individual belief states are consistent, so coherence fails.

This leaves intersection as the last aggregation method standing – the only one that has all of the above desirable features. But many will find intersection to be an unacceptable aggregation procedure, since groups will be forced to withhold on any matter for which there is any disagreement. But epistemic aggregation should allow groups to have opinions and act on those opinions even when there is some disagreement. Can we do better?

3. MOTIVATING CREDENCES (AND MORE)

One response to the problem of aggregating outright beliefs is to try to aggregate some other sort of epistemic state instead. Outright beliefs leave us with only a few options for characterizing a group’s view on any particular matter (belief, disbelief, suspension).

5 Locality is an intuitively plausible principle, and it makes it much easier to compute group beliefs. But we also judge the tree by its fruits, and we find that there is a way of satisfying locality that we prefer to any known way of violating locality. See Douven and Romeijn (2007) for a technical analysis of locality, and Douven and Williamson (2006) for more on the difficulties of weakening locality.

Maybe we'll find greener pastures for epistemic aggregation if we aggregate credences (also sometimes called degrees of confidence) rather than outright beliefs.

For our purposes, we will assume that the traditional Bayesian picture of credences is the right one. The traditional Bayesian picture says that there are continuum many possible epistemic states, which together have the same structure as the real numbers between 0 and 1. Like the real numbers in the unit interval, there are continuum many degrees of confidence with absolute certainty of truth (credence 1) at the top and absolute certainty of falsehood (credence 0) at the bottom. Like the real numbers, credences can be added, multiplied, divided and so on to arrive at other credences.⁶ Because credences have the same structure as the unit interval, we can represent them as a function from propositions to the reals between 0 and 1.

Traditional Bayesians impose two kinds of coherence requirements on rational credences: a synchronic constraint (which applies at each time) and a diachronic constraint (which applies across different times). The synchronic requirement is probabilistic coherence, the requirement that the agent's credences at any time should form a probability function. That is, an agent's degrees of belief must (at a minimum) conform to the axioms of the probability calculus:

- (1) $\Pr(p) \geq 0$, for any proposition p .
- (2) $\Pr(t) = 1$, for any tautology t .
- (3) $\Pr(p \vee q) = \Pr(p) + \Pr(q)$, for any inconsistent propositions p and q .

The diachronic requirement is **conditionalization**. This constraint relies on a notion of conditional probability, which we can define given the axioms above. The probability of p conditional on q is equal to the probability of p and q divided by the probability of q (assuming that $\Pr(q) > 0$).

$$\Pr(p | q) =_{df} \frac{\Pr(p \wedge q)}{\Pr(q)}.$$

Conditionalization mandates that when an agent has evidence E at time a and evidence $E+$ at time b (where $E+$ implies E), then the probability of any proposition p at b should equal the conditional probability of p given $E+$ at a . That is, $\Pr_b(p) = \Pr_a(p|E+)$.

Putting this all together, the basic story of Bayesian epistemology goes something like this. Rational agents assign non-negative credence to various possible worlds (possibilities that are disjoint and exhaustive) in such a way that the sum of those credences is 1. When rational agents gain evidence that is inconsistent with some of those possible worlds, they adjust their credences in those worlds to 0, and proportionally increase their credences in the remaining worlds so that the sum of their credences returns to 1.

By moving from outright beliefs to degrees of belief, we give epistemic states more structure. One might think that by doing so we can get enough room to solve some of the problems of epistemic aggregation. In fact, we can. One might think that by doing

⁶ So long as these operations make sense – that is, so long as the operations do not take us outside of the unit interval. Credence 0.4 can be added to credence 0.5, but there is no credence that is the result of adding credence 0.4 and credence 0.8.

so we can get enough room to solve all the problems of epistemic aggregation. In fact, we cannot. But maybe that means that even degrees of belief do not give us enough structure. After seeing what can and cannot be done within a credal framework, we will explore an alternative that gives us a new option for even more options for how to aggregate epistemic states.

4. POSSIBILITIES AND IMPOSSIBILITIES IN CREDAL AGGREGATION

We can do a lot with credal aggregation, just not everything we want.

4.1 Possibility

Let us start with the things we can do. We saw earlier that intersection is the only procedure for aggregating outright beliefs that satisfies coherence, locality, anonymity and unanimity. To see whether or not there are aggregation procedures for credences that have all four desirable features, we need to first describe what those features should look like in the case of aggregating credences.

In the context of credences, **coherence** says that the aggregation of coherent credences always produces coherent credences. What are coherent credences? Based on the traditional Bayesian picture that we are assuming, an agent has coherent credences just in case her credences are probabilistically coherent and updated by conditionalization. So coherence for credences has two parts: **probablism**, the requirement that probabilistically coherent credence functions be aggregated into a probabilistically coherent credence function, and **conditionalization**, the requirement that individual credences that obey conditionalization are aggregated into group credences that obey conditionalization.

You will recall that **locality** is the idea that what a group thinks about a proposition should be fully determined by what the individuals think about that proposition. So when aggregating credences, we should be able to determine how confident the group is about p just by looking at how confident the individuals are about p . More precisely, locality requires there to be a function $f: [0, 1]^n \rightarrow [0, 1]$ such that for any sequence of individual credence functions $\langle C^1, C^2, C^3 \dots \rangle$, any aggregate C^* of that sequence of credence functions, and any proposition p , $C^*(p) = f\langle C^1(p), C^2(p), C^3(p) \dots \rangle$.⁷

As in the case of belief, we will be thinking of credence aggregation as a function that takes sequences of credence functions as inputs and generates a group credence. **Anonymity** then works precisely the same way that it did with the aggregation of beliefs. A method of aggregation satisfies anonymity just in case permuting a sequence S and aggregating the result is equivalent to aggregating the original sequence S .

Unanimity is just what one would expect: if everyone in a group has credence x that p , then the group also has credence x that p .

We now want to note a few important implications of these constraints. Probablism and locality together entail unanimity (the proof is not hard and so we leave its reconstruction to the reader). A more substantial result is that arithmetic averaging is the only

7 Locality for credences goes by a number of different names in the literature, including ‘strong label neutrality’, the ‘strong set wise function property’ and ‘the context-free assumption.’ These are all too clumsy and long-winded for our tastes.

method of credal aggregation that satisfies both probabilism and locality.⁸ And the only method that satisfies probabilism, locality and anonymity is straight arithmetic averaging.

4.2 Impossibility

Tragically, straight arithmetic averaging has a grievous flaw: it produces group credences that violate conditionalization even when all individuals in the group obey conditionalization. Suppose that we have three worlds and two agents whose credences in those worlds are represented by the top two rows in the following table:

	w_1	w_2	w_3
C^1	1/4	1/2	1/4
C^2	1/2	1/4	1/4
C^*	3/8	3/8	1/4

Straight averaging gives us the group credences in the bottom row. Now suppose that both individuals learn that w_1 is false and conditionalize on that fact. That gives us updated results in the top two rows:

	w_1	w_2	w_3
C_+^1	○	2/3	1/3
C_+^2	○	1/2	1/2
C_+^*	○	7/12	5/12

If we then use straight averaging to get the group credences, the group has the credences in the bottom row. But notice that the new group credence C_+^* is *not* the result of conditionalizing C^* on the group’s new evidence that w_1 is false. If it were, we would have $C_+^*(w_2) = 3/5$ and $C_+^*(w_3) = 2/5$. So straight averaging fails to satisfy conditionalization.⁹

8 See McConway (1981) and Wagner (1982).

9 There is an aggregation procedure that satisfies anonymity and conditionalization but not locality: taking the renormalized geometric mean of elementary possibilities. Elementary possibilities are the most specific possibilities within an epistemic framework. They may be thought to be some objectively definable set of epistemically possible worlds, or they may be defined framework by framework. The geometric mean of n numbers is the n^{th} root of the product of those numbers. If we take the geometric mean of the individuals’ credences in each world and add them together, that sum will be less than 1 (if the individuals’ credences are not all the same.) Hence it is important to renormalize, dividing the geometric mean of each world by the sum of all the geometric means.

We do not like taking the renormalized geometric mean of elementary possibilities, for a number of reasons. First, we *really* like locality. Locality makes sense, and it makes aggregation a lot easier. Secondly, the credence that the renormalized geometric mean of elementary possibilities ends up assigning to a possibility will depend on which possibilities are taken to be elementary. Thus, unless there is an objective set of epistemically possible worlds to use, the contingency that results is deeply unpalatable. Thirdly, the renormalized mean is unwieldy: if any individual has credence ○ in a proposition, the group must have credence ○ in that proposition. (Similarly, an individual can force the group to have arbitrarily low credence in a proposition by having a sufficiently low credence in that proposition.) This is particularly bad news if the renormalized geometric mean of elementary possibilities uses an objective set of

5. MOTIVATING CREDAL PAIRS

We can aggregate individual credences at a time into group credences at that time just the way we want. Straight averaging (and only straight averaging) obeys probabilism, locality, anonymity and non-triviality. But we cannot keep all of these features and still have group credences evolve over time the way we want them to. Straight averaging violates conditionalization.

It is not entirely surprising that these group credences do not evolve the way we want them to. Credences change due to evidence, but our method for aggregating credences does not – and cannot – take evidence into account. Credal aggregation applies to probability functions, and probability functions underdetermine evidential states. Any time an agent has a probability function because he conditionalized on some evidence, some other agent could have the same probability function without any evidence at all.

Perhaps credal aggregation is too restrictive. It pays attention only to unconditional credences at a given time; it is blind to how those unconditional credences came about. Bayesian epistemology has more structure than credal aggregation pays attention to. Unconditional credences are not primitive; they result from conditionalizing a prior on evidence. To capture this structure, let us represent an agent as having a **credal pair** – an ordered pair $\langle P, E \rangle$ of a prior probability and an evidential state.¹⁰

We can use conditionalization to calculate the unconditional credences for any credal pair, but we will know more about where those unconditional credences came from. Credal pairs contain strictly more information than unconditional credences. The additional structure of credal pairs presents new possibilities for epistemic aggregation.

6. AGGREGATION PROCEDURES FOR CREDAL PAIRS

How shall we derive a group credence function from a set of credal pairs?¹¹ Two obvious options present themselves: (1.) Use each individual's credal pair to calculate that individual's credence function, and then aggregate those individual credence functions into a group credence function. (2.) Aggregate the individuals' credal pairs into a group credal pair (aggregating the individuals' priors into a group prior and the individuals' evidence into a group evidence), and then use that group credal pair to calculate a group credal function. We'll call Option 1 *Calculate, then Aggregate* and Option 2 *Aggregate, then Calculate*.

epistemically possible worlds. If there are uncountably many epistemically possible worlds, every person must give credence \circ to nearly all of the worlds. And if every elementary possibility is assigned credence \circ by some individual, then the group assigns \circ to every world, thus giving us a violation of probabilism.

For a defense of taking the renormalized geometric mean, see Russell, Hawthorne, and Buchak (manuscript).


- 10 We find it most natural to think of the priors in credal pairs as ur-priors, prior probabilities that represent completely aprioristic beliefs. An agent's ur-prior is completely independent of his evidence, so credal pairs using ur-priors cleanly separate evidential considerations from non-evidential considerations. Unconditional credences from other times can be used as priors, but which time's unconditional credences are chosen can affect the result of aggregation.
- 11 We are not necessarily aggregating individuals' credal pairs into a group credal pair. Nonetheless, one of the two aggregation methods we will consider can be used to produce a group credal pair.

Calculate, then Aggregate is equivalent to averaging group members' credences, a procedure we have already explored.¹² *Aggregate, then Calculate* is a new procedure that cannot be formulated given only group members' credences. At a glance, the two procedures seem like comparably reasonable ways to aggregate credal pairs. Let us look at them in more detail.

6.1 *Calculate, then aggregate (straight averaging)*

Calculate, then Aggregate does to credal pairs exactly what straight averaging does to probability functions. *Calculate, then Aggregate* is just a relabeling of straight averaging. Although *Calculate, then Aggregate* nominally employs more structure than straight averaging does, that additional structure does no work.

If a group uses *Calculate, then Aggregate*, the credal pairs of the group members at a given time will determine the credence functions of the group members at that time, which in turn will determine the group credence function at that time.¹³ If the group gains evidence, its aggregate credence function will not be updated directly. The aggregate credence function will change only because the individual credal pairs will change, which in turn will change the individual credence functions. At any time, the values assigned by the individual credence functions will determine the values assigned by the group credence function – each member of a group of size n will contribute $1/n$ of his credence in a proposition to the group's credence in that proposition.



6.2 *Aggregate, then calculate*


Aggregate, then Calculate does something with credal pairs that cannot be done with probability functions. It uses the additional structure of credal pairs to respond to the two factors that underlie an agent's probability function: prior probability and evidence.

If a group uses *Aggregate, then Calculate*, the credal pairs of the group members at a given time will determine the credal pair of the group at that time, which in turn will determine the group credence function at that time. The individuals' priors will be aggregated into a group prior, and the individuals' evidence will be aggregated into the group's evidence. The aggregation of the priors will be a straight average (thus satisfying coherence, locality, anonymity and unanimity), and what the aggregation of the evidence will be

12 Assuming that the method of credal aggregation is a straight average. All other methods of credal aggregation remain possible, but a straight average does better at satisfying the desiderata for credal aggregation procedures than the alternatives.

13 We assume that the individuals update their credences by conditionalization.

should depend on how the evidence is understood. The group credence function will change only because the group’s evidence will change; the group’s prior will be constant. Each member of a group of size n will contribute $1/n$ of his prior in a proposition to the group’s prior in that proposition.



6.3 Weighting members’ credences by epistemic success

Calculate, then Aggregate gives each individual’s credence function an equal effect on the group credence function. *Aggregate, then Calculate* gives each individual’s prior an equal effect on the group prior. We do not think one of these two sorts of equality is obviously superior to the other. Each method of epistemic aggregation seems to give fair treatment to each individual in the group.

Interestingly, although *Aggregate, then Calculate* averages the individuals’ priors, it can nonetheless be viewed as averaging the individuals’ credence functions, just like *Calculate, then Aggregate* does.¹⁴ The difference is that while *Calculate, then Aggregate* provides a straight average of the individuals’ credences, *Aggregate, then Calculate* provides a weighted average of the individuals’ credences – a dynamic weighting that is determined at each time by how much of the individual’s prior remains unfalsified by evidence at that time. Let us say that the amount of an individual’s prior that remains unfalsified by evidence determines that individual’s success at predicting the content of that evidence.

Given *Aggregate, then Calculate*, the more successful an individual was at predicting the content of the evidence, the more effect that individual’s credences will have on the group’s credences. One can thus think of *Calculate, then Aggregate* as providing an unweighted average of the individuals’ credences, and *Aggregate, then Calculate* as providing an average of the individuals’ credences that is weighted according to epistemic success.

Proof: Let C^* be the aggregate prior, the C^i the individual priors, C_e^* the result of conditionalizing the aggregate prior on e , and the C_e^i the result of conditionalizing the individual priors on e . Furthermore, say that $success_i(e) = C^i(e) / \sum C^i(e)$ is the **relative success** of i at predicting e . We want to show that if the group prior is the result of averaging the individual priors, then the unconditional group credence $C_e^*(p)$ is equivalent to the sum of the individual unconditional credences $C_e^i(p)$ weighted by the relative epistemic success of

¹⁴ At least when all the individuals and the group as a whole have the same evidence. For the purposes of this section, we make this assumption.

each individual i at predicting e . Let i range over the individuals in a finite group of size n . Then:

$$\begin{aligned}
 C_e^*(p) &= \frac{C^*(p \wedge e)}{C^*(e)} \\
 &= \frac{(\sum C^i(p \wedge e))/n}{(\sum C^i(e))/n} \\
 &= \frac{\sum C^i(p \wedge e)}{\sum C^i(e)} \\
 &= \sum \left(\frac{C^i(p \wedge e)}{\sum C^i(e)} \right) \\
 &= \sum \left(\frac{C^i(e)}{\sum C^i(e)} \cdot \frac{C^i(p \wedge e)}{C^i(e)} \right) \\
 &= \sum (\text{success}_i(e) \cdot C_e^i(p))
 \end{aligned}$$

This particular notion of epistemic success is admittedly somewhat peculiar (there is no guarantee that agents who invested more credence in falsified possibilities are worse at distributing their credences among unfalsified possibilities), but it is a notion of epistemic success that can be defined objectively.¹⁵ Every individual gets an equal opportunity to contribute to the group's credences.

It is not unreasonable to think that each individual's credences should contribute equally to the group's credences, but neither is it unreasonable to think that individuals who have had more predictive success (those who gave less credence to falsified hypotheses) should contribute more to the group's credences than individuals with less predictive success (those who gave more credence to falsified hypotheses). A bad track record might not mean anything, but it also might indicate unreliability or unreasonableness. Consider the following case:

Dissimilar Angels

Two angels, Wise and Foolish, come into existence at the start of creation. Wise and Foolish have very different opinions about whether or not God will turn the universe purple on any particular day. For each day n , Wise has credence $1/2n$ that God will turn the universe purple on that day if God has not done so already. Foolish, on the other hand, has credence $1 - (1/2n)$ that God will turn the universe purple on day n if he has not done so thus far. As time goes on and God predictably does not turn the universe purple, Foolish will become arbitrarily confident that *today is the day* God will finally turn the universe purple and Wise will become arbitrarily confident that God will not turn the universe purple today.

If Wise and Foolish aggregate their credences using straight averaging, they will have a group credence of $1/2$ that God will not turn the universe purple on any particular day, given that he has not yet done so.¹⁶

¹⁵ That is, without prejudging any individuals or any possibilities.

¹⁶ The renormalized geometric mean also gives this verdict in this case.

Weighting by success, on the other hand, gives very different results. On each of the first two days, Wise and Foolish will have a group credence of $1/2$ that God will turn the universe purple. But by the third day, the relative success of the two angels will start having an effect on the weight their credences receive when calculating the group credences. Foolish will get a weighting of $1/4$ and Wise will get a weighting of $3/4$, resulting in a group credence of $5/16$ that God will turn the universe purple. As time goes on and Wise continues to make good predictions, her opinion will continue to count for more, resulting in group credences that are arbitrarily close to her own.

7. THE VIRTUES OF WEIGHTING BY SUCCESS

In the last section, we showed that *Aggregate, then Calculate* is equivalent to weighting the opinions of the individuals in a group by their success at predicting the evidence. The straight averaging of priors ensures that agents all count equally when it comes to establishing the group priors – everyone has an equal opportunity at having their future opinions influence the group’s opinion. But as the evidence starts coming in, those with more success at predicting the evidence will have greater influence on the group credences and those with less success will have less.

In previous sections we discussed four virtues for aggregation procedures – locality, anonymity, coherence and unanimity. How should we understand these virtues when it comes to credal pairs?

It should be clear by now what **anonymity** is going to say: merely permuting a sequence of credal pairs (rearranging which individual has which credal pair) never changes which credal pair a group has.

The most natural extension of **locality** to credal pairs says that we can determine both the group prior and the group evidence point-wise. Representing priors as probability functions and evidence as a characteristic function, we can state the requirement like this. There are functions $f: \{0,1\}^n \rightarrow \{0,1\}$ and $g: \{0,1\}^n \rightarrow \{0,1\}$ such that $\langle P^*, E^* \rangle$ is an aggregate credal pair for the sequence $\langle P^1, E^1 \rangle, \langle P^2, E^2 \rangle, \langle P^3, E^3 \rangle \dots$ just in case for all propositions p , both $P^*(p) = f\langle P^1(p), P^2(p), P^3(p) \dots \rangle$ and $E^*(p) = g\langle E^1(p), E^2(p), E^3 \dots \rangle$.

What about coherence? Two things should almost go without saying. Rational priors are probabilistically coherent and rational evidence sets are logically consistent. So **coherence** for credal pairs should at the very least require that when we aggregate pairs of probabilistically coherent priors and logically consistent evidential states, we should get back a credal pair with a probabilistically coherent prior and a logically consistent evidential set. But there is no need for the priors in credal pairs to obey conditionalization. The prior in a credal pair represents what an individual thinks before evidence comes in, not what the agent thinks in light of that evidence. The agent’s unconditional credences should obey conditionalization, which they will if the agent’s priors remain fixed across time. We think this is good reason to require priors to remain fixed. Diachronic coherence then requires the group’s priors to remain fixed whenever the individual priors remain fixed.

Unanimity is also easily defined. If every individual’s credal pair has a prior that assigns x to p , then the group’s credal pair has a prior that assigns x to p , and if everyone in the group has some proposition as part of their evidence, then the group has that proposition as part of its evidence.

It is not hard to see, given what we have said already, that the only aggregation method for credal pairs that meets all these criteria is the one that we arrive at by taking the straight average of individual priors and the intersection of individual evidential sets. The proof that the aggregation of priors must be straight averaging is the same as it was in the case of credences, and the proof that evidence aggregation must be intersection is the same as it was in the case of belief.

This is worth saying again. When it comes to credal pairs, there is precisely one aggregation method that achieves the ideals of anonymity, locality, coherence and unanimity: averaging priors and intersecting evidence.

Thinking of credences as states that are calculated from credal pairs, we can now also explain why anonymity and locality should be violated if credences are erroneously viewed as fundamental. Anonymity says that switching who thinks what should never change what the group thinks as a whole. But when we are weighting credences by the relative epistemic success of those who have them, we should expect changes in who thinks what to change the credences of the group as a whole.

A similar line of reasoning explains why locality for credences is no good. Locality says that when two groups have the same pattern of credences in a proposition, the groups should have the same credences in that proposition. But if we are weighting epistemic success, then the track records of individuals also matter. So there is no reason to expect the distribution of individual credences in a proposition to fully fix the group credences in that proposition.¹⁷

Straight averaging priors and intersecting evidence is the only method for aggregating credal pairs that satisfies coherence, locality, anonymity and unanimity. If we take this method for aggregating credal pairs and view it from the more limited perspective of probability functions, it violates locality and anonymity. The import of the extra structure in credal pairs explains why no method for aggregating mere probability functions can satisfy coherence, locality, anonymity and unanimity.

We should note one thing about evidence aggregation. The proof that the procedure for evidence aggregation must be intersection (in order to satisfy the four desiderata in the belief section) depends on the possibility of agents having evidence that is collectively inconsistent. But if evidence is factive, then no such thing is possible. It would then be much easier to find an aggregation procedure that satisfies coherence. Thus, those who are willing to commit to the factivity of evidence will have more options for aggregating credal pairs. The straight averaging of priors is still required in order to satisfy locality, probabilism and anonymity, but there will be more latitude to combine the straight averaging of priors with other methods of evidence aggregation.

Because it makes use of the extra structure encoded in credal pairs, *Aggregate, then Calculate* can possess all four of the important virtues for aggregation procedures. As such, we judge this to be the best aggregation procedure for credal pairs.

¹⁷ While locality fails for the aggregation of unconditional credences, a weaker **neutrality** condition holds for *Weighting by Success*. The weaker neutrality condition says that if we fix the priors and evidence for a particular group, then there is a function $f: [0,1]^n \rightarrow [0,1]$ such that for any sequence of unconditional individual credence functions $\langle C^1, C^2, C^3 \dots \rangle$ such that for any proposition p , $C^k(p) = f_{i,g}(C^i(p), C^2(p), C^3(p) \dots)$. In other words, an aggregation method satisfies neutrality if whenever the individuals in a group have the same unconditional credences in two propositions p and q , then the group has the same credence in p and q . This gets at what may be one of the motivating ideas behind locality: that credence aggregation should be topic neutral. Locality entails neutrality, but clearly not *visa versa*.

8. CONCLUSION

We like *Aggregate, then Calculate*. It captures everything we wanted from an epistemic aggregation procedure. But maybe we are missing something. We are open to the possibility that some other aggregation procedure is even better. We are not as attached to any particular theory of epistemic aggregation as we are to an overall method: don't get pushed around by impossibility results. Impossibility results do not show the limitations of social epistemology; they show the limitations of particular frameworks for doing social epistemology. We see no reason not to be confident that any reasonable desiderata can be accommodated within a reasonable epistemological framework – that framework may just need a bit more structure than other epistemological frameworks. Impossibility results show that we cannot get everything we always wanted the way we always expected, but we are hopeful that through the unexpected, we will get everything we always wanted anyway.

REFERENCES

- Douven, Igor, and Williamson, Timothy. 2006. "Generalizing the Lottery Paradox." *British Journal for the Philosophy of Science*, 57(4): 755–79.
- Douven, Igor, and Romeijn, Jan-Willem. 2007. "The Discursive Dilemma as a Lottery Paradox." *Economics and Philosophy*, 23: 301–19.
- Fitelson, Branden, and Easwaran, Kenny. 2013. "Accuracy, Coherence, and Evidence." Paper presented at Soberfest, Madison, WI, 24–25 May.
- List, Christian, and Pettit, Philip. 2004. "Aggregating Sets of Judgments: Two Impossibility Results Compared." *Synthese*, 140: 207–35.
- McConway, Kevin. 1981. "Marginalization and Linear Opinion Pools." *Journal of American Statistics*, 76: 410–14.
- Russell, Jeffrey Sanford, Hawthorne, John, and Buchak, Lara. "Groupthink." Manuscript.
- Wagner, Carl. 1982. "Allocation, Lehrer Models, and the Consensus of Probabilities," *Theory and Decision*, 14: 207–20.

DANIEL BERNTSON is a Ph.D. candidate at Princeton University with a Masters from Brown University. He works primarily in metaphysics, epistemology, and the philosophy of science.

YOOAV ISAACS is a Ph.D. student at Princeton University and has Masters degrees from the University of Chicago and the University of Oxford. His main interests are in epistemology, decision theory, philosophy of religion and ethics.
