Group Credences

Belief and its Limits (BaiL); Seminar 6 ANU
Instructor: Joshua Pearson; joshuaedwardpearson@gmail.com 11/11/25

I. Opening Task: Dutch Books and Probabilism

Suppose you know two facts about Steve:

1. He'll always choose the option he thinks will, on average, produce the most money.

   So, where $c(\cdot)$ represents Steve's credences, he'll buy a bet that promises a profit of $\$x$ if p and a loss of y if not-p iff $(c(p) \times x) + (c(\neg p) \times -y) \geq 0$.

   And he'll sell such a bet iff $(c(p) \times x) + (c(\neg p) \times -y) \leq 0$.

2. He's $\frac{1}{3}$ confident a die I'll roll tomorrow will land on 1 or 2, $\frac{1}{3}$ confident it will land on 3 or 4, but only $\frac{1}{2}$ confident it will land on 1 or 2 or 3 or 4. (He violates "**Finite Additivity**")

That is, he always tries to "maximise expected utility" (the orthodox view in decision theory), but cares only about money.

Two notes:

- Also assume that Steve satisfies "**Negation**": $c(p) = 1 - c(\neg p)$.
- **TIP:** given this assumption, it follows that Steve will buy a bet on p that pays \$1 so long as it costs no more than $\$c(p)$, and will sell such a bet so long as it costs at least $\$c(p)$.

  See Hedden (2011) for useful discussion about how dutch-book arguments relate to orthodox decision theory.

You can offer Steve a series of bets to either buy or sell that is guaranteed to make you money off of him (a "dutch book")—can you think of an example?

Do you think this shows Steve is epistemically irrational?

II. Group Credences

Last week we looked at group *full* belief. It's thus natural to move to group credences, like we moved from individual belief to credences. But something interesting happens in the literature here. Note that the following two questions aren't obviously equivalent:

1. Under what conditions does a group have credence x in p, or under what conditions is such a credence rational?

2. What is the best way for a group to aggregate credences of their members so as to coordinate group action?

If Summativism is correct (group beliefs just summaries of its members beliefs) question 2 may inform question 1. But:

- *Summativism might not be correct!* A non-summativist might see value in question 2, but think it's importantly distinct from 1.

- *Even if summativism is correct*, the best way for a group to aggregate credences so as to coordinate group action might be distinct from the credence the group should actually be interpreted as having.

Nonetheless, the literature seems very focused on question 2.[1] **We'll follow the literature on this point**—just note that there may be importantly different questions we can ask here.

[1] Indeed, Russell, Hawthorne and Buchak are quite clear about this in the abstract of their paper: "How should a group with different opinions (but the same values) make decisions?"

III. Averaging and Dutch Books

III.1—An Average Proposal

So, how should a group aggregate credences so as to coordinate action? A natural proposal is that the group should take the average of the credences of its members.[2] Let's introduce some simplifying assumptions and terminology to get the proposal on the ground.

[2] Or perhaps a weighted average, if for some reason not every member should be given equal voice.
See also Moss (2011) for an interesting alternative picture, under which this average also needs to take into account the different ways members of a group might epistemically value different credences.

- We'll assume we're interested in one group with n members.
- Let $c_i(\cdot)$ denote a possible credence function of member i.
- Let C denote a possible assignment (a sequence) of credence functions to each member:

$$C = \langle c_1(\cdot), c_2(\cdot), ..., c_n(\cdot)\rangle$$

- Let $C(A)$ denote the sequence of each members credence in A according to C:

$$C(A) = \langle c_1(A), ..., c_n(A)\rangle$$

- Let $ag(\cdot)$ denote the desired aggregation function, that takes us from a sequence of individual credence functions to the group credence function. $ag(C)$ thus the credence function ag assigns to C, and $ag(C)(A)$ is the corresponding aggregted credence in A.

We can state the average proposal as follows:

(Averaging) For all propositions A and assignments of credence functions $C = \langle c_1(\cdot), ..., c_n(\cdot)\rangle$:

$$ag(C)(A) = \frac{c_1(A) + ... + c_n(A)}{n}$$

A good feature: the group credence will always be synchronically probabilistically coherent whenever the individual credences are probabilistically coherent.

A potential bad feature: Russell, Hawthorne and Buchak—"RHB"—illustrate if a group's credences conform to **Averaging**, they will be susceptible to a "diachronic dutch book"...

III.2. A long story about Acme...

Acme Corp Part 1—Anvils & Balloons Acme are divided over whether to invest into a anvil factory. If the factory is unsuccessful, they'll lose $11k—the amount it costs to invest. If successful, they'll gain $21k—resulting in a $10k profit. While they all agree they should invest iff doing so yields more money on average than not investing,[3] they disagree on how likely the investment is to succeed.

Half of them—the "avilites"—think the investment is $\frac{2}{3}$rds likely to succeed. So, for them, the average profit of investing is $(\$10k)(\frac{2}{3}) + (-\$11k)(\frac{1}{3}) \approx \$3k$ profit. So, they are pro-investment.[4]

The other half only think the investment is $\frac{1}{3}$rd likely to succeed. So, for them, the average profit is $(\$10k)(\frac{1}{3}) + (-\$11k)(\frac{2}{3}) \approx -\$4k$. They are anti-investment.

Acme decide to act according using credences obtained by **Averaging**. On average, they are $\frac{1}{2}$confident investment will be successful. According to this probability, investment will on average yield $(\$10k)(\frac{1}{2}) + (-\$11k)(\frac{1}{2}) \approx -\$0.5k$. **They decide not to invest.**

The situation with an investment into a balloon factory is similar, but reversed. Now the anvilites are only $\frac{1}{3}$ confident the investment will succeed, and the other half—the "balloonites"—are $\frac{2}{3}$rds confident of success. But the costs and rewards are the same, and Acme again act on their average credence—meaning they decide not to invest.

Acme Corp Part 2 — A Bet. Acme are approached by a cunning stockbroker offering them a bet.

Bet 1: For a price of $20k, Acme will win $37k ($17k profit) if exactly one of the anvil or balloon factories succeeds, and nothing otherwise.

Everyone in Acme agrees that the factories operate independently from one another. This means the both the anvilities and the balloonites are $\frac{2}{9}$ confident both factories will succeed and $\frac{2}{9}$ confident both will fail.[5] Further, both the anvilities and the balloonites are $\frac{5}{9}$ths confident—albeit for different reasons[6]—that exactly one factory will succeed. On average, then, the bet pays $\$(17k)(\frac{5}{9}) + (-\$20k)(\frac{4}{9}) \approx \$0.560$. They agree to take the bet!

Acme Corp Part Three — A problem. A year goes by an anvils do badly, but the fate of the balloons is undecided. The stockbroker now offers them a second bet:

Bet 2: for a price of $18k, they will win $37k ($19k profit) if the balloon factory fails.

Again, taking the averages of their credences, Acme take this bet to on average yield $500 bucks, so they take it.

But now something weird has happened. Suppose the balloon factory succeeds. Then Acme will get gain $17k from bet 1, and lose $18k from bet 2. And suppose the balloon facotry fails. Then Acme will lose $20k from bet 1 and gain $19k from bet 2. Eitherway, they are *guaranteed a loss* (and the stockbroker a guaranteed profit)![7] Even worse...

[3] That is, they all conform the orthodox decision-theoretic rule that one should "maximize expected utility", mentioned for the opening task. (And they also only care about money....)

[4] On the other hand, not investing will return a profit of $0; so they take investing to be the on-average better choice.

[5] This is because (remember from class 4!) when we have probabilistic independence between p and q, $Pr(p\&q) = Pr(p) \times Pr(q)$

[6] For the anvilities, it's because the probabiity of anvil-success without balloon-success is $(\frac{2}{3})(\frac{2}{3})$, and the probability of balloon-success without anvil-success is $(\frac{1}{3})(\frac{1}{3})$, giving a total of $\frac{5}{9}$. The balloonites are symmetrically more confident of balloon-success than anvil-success, meaning the numbers work out the same.

[7] And they haven't been "tricked" by the stockbroker: she presented the bets honestly, and Acme do indeed take these bets to be fair by their own lights... (Of course, there's perhaps another kind of trickery going on—the stockbroker is exploiting this company—but at least she's not tampering with Acme's agency in any sense.)

Acme Corp Part 4—the "Dutch Book". If the anvil factory had done well, the stockbroker would have *instead* offered them:

Bet 3: Pay $18k, win $37k ($19k profit) if the balloons do well, nothing otherwise.

Acme would have thought this bet was fair, too—yielding on average $500.

But again, in this counterfactual case, if the balloon factory succeeded, Acme lose $20k from the first and gain $19k from the third, and if the balloon factory failed, they gain $17k from the first bet but lose $18k from the third. Acme are against guaranteed to lose money, and the stockbroker guaranteed to win.

II.2. ...And?

Of course, the above story is extremely specific. That **Averaging** can lead to a sure loss in this specific scenario might not be so bad.

The problem is that this is a symptom of the following fact: if a group credence obeys **Averaging**, then they will not update by conditionalization: where $m > n$, and e_m is the evidence gained between times t_n and t_m, $G^{t_m}(p)$ will not in general be equal to $G^{t_n}(p \mid e_m)$.[8]

[8] Here's a simple example. Suppose our group is made of two people, with credence functions c_1 and c_2. They are both certain that a fair coin is either $\frac{9}{10}$ biased towards heads or biased towards tails, but their credences in these hypotheses differ as follows: $c_1(HeadsBiased) = c_2(TailsBiased) = \frac{9}{10}$.

- Their average credence in *HeadsBiased* is $\frac{1}{2}$.
- Using Bayes' Theorem, on seeing the coin land heads:
  - $c_1(HeadsBiased \mid Heads) = \frac{(\frac{9}{10})(\frac{9}{10})}{(\frac{9}{10})(\frac{9}{10})+(\frac{1}{10})(\frac{1}{10})} = \frac{81}{82}$
  - $c_2(HeadsBiased \mid Heads) = \frac{(\frac{1}{10})(\frac{1}{10})}{(\frac{1}{10})(\frac{1}{10})+(\frac{9}{10})(\frac{9}{10})} = \frac{1}{82}$
  - So, if they update by conditionalisation, the new group credence will be the midpoint of these: $\frac{41}{82} = \frac{1}{2}$
  - Yet this is not the same result as if the group credence had updated by conditionalization: $G(HeadsBiased \mid Heads) = \frac{(\frac{9}{10})(\frac{1}{2})}{(\frac{9}{10})(\frac{1}{2})+(\frac{1}{10})(\frac{1}{2})} = \frac{9}{10}$

The above story is just one instance of a result from Lewis (2011) that agents who violate conditionalization will be susceptible to these "diachronic dutch books". Lewis further shows that credences which *do* update by conditionalization will *not* be susceptible to them.

Let's turn all of this into an argument against **Averaging**:

1. A group credence that satisfies **Averaging** can be diachronically dutch-booked.
2. If a group is susceptible to diachronic dutch books, it is irrational.
3. Group credences should not be determined by **Averaging**.

...Are we convinced by this argument against **Averaging**?

III. Can Groups Satisfy Conditinalization?

Let's suppose the above has convinced us that groups should not determine their credences according to **Averaging**. If so, we should think that group credences should be aggregating in a way consistent with conditionazation. Let's define this constraint formally.

- Let $c_i \mid A$ be the result of conditionalisaing member i's credence on A, and let $C \mid A$ be the sequence of the credences in C updated by A: $\langle c_1 \mid A, ..., c_n \mid A \rangle$.

Conditionalization: $ag(C \mid A) = ag(C) \mid A$.

In words: *Aggregating the member's credences conditionalized on A is equal to conditionalizing the group's prior aggregate credence on A.*

RHB argue that satisfying conditionalization is not easy—it requires we give up other plausible principles of aggregation.

Here are two such principles

Non-Dictatorship There is no i such that, for every possible sequence of credences C, $ag(C) = c_i$.

In words: *There should be no member i of our group such that, regardless of what everybody else's credences are, the aggregate credence is just the credences of i!*

Irrelevant Alternatives For any two sequences of credence functions C and C', if $C_i(A) = C'_i(A)$ for each i, then $ag(C)(A) = ag(C')(A)$.

In words: *determining a group's credence in A does not require you to look at their credence in other propositions.*

Fact 1: No rule for aggregation satisfies all of Conditionalization, Non-Dictatorship and Irrelevant Alternatives![9]

[9] Any thoughts on which principle might be the weak link here?

And here are three more such principles:

Anonymity If C' is a perumtation of a sequence C, then $ag(C) = ag(C')$

In words: *A group's credences should not depend on which particular members believe which particular claims—they can be "anonymous".*[10]

[10] Anonymity entails Non-dictatorship.

Unanimity If $C = \langle c_1 \cdot), ..., c_n(\cdot)\rangle$ and $c_1(A) = c_2(A) = ... = c_n(A) = x$, then $ag(C)(A) = x$

In words: *If the group all agree that A is x likely, then the aggregated credence in x should be A.*

The next one is the most complex. But the idea behind it is intuitive: we shouldn't need to know the content of particular propositions—only their names "A", "B"—to aggregate the group's credences.

Let π be a permutation of worlds in W. And let $\pi[c_i(w_1)] = c_i(\pi(w))$.

Neutrality For any possible sequence of credences $C = \langle c_1(\cdot), ..., c_n(\cdot)\rangle$ and "world permutation π, $ag(\langle \pi[c_1(w)], ..., \pi[c_n(w)]\rangle) = \pi[ag(C)(w)]$.

Fact 2: No rule for aggregation satisfies all of Conditionalization, Anonymity, Unanimity and Neutrality![11]

[11] Again... which principle should we give up?

IV. Alternatives to Averaging?

IV.1 The Fixed Prior Rule

(Fixed Prior Rule) Let Pr be a specified probability function, let C be a sequence of credence functions, an let E be the conjunction of propositions A such that for all i, $C_i(A) = 1$. Then:

$$ag(C) = Pr \mid E$$

The idea: agree on some credence function at the outset. Then just conditionalize on this function.[12]

[12] We could also make the relevant notion of evidence a "pooling" rather than "common ground" notion: E is the conjunction of propositions A such that for *some* i, $c_i(A) = 0$. But this can't be applicable in every case: what is every proposition such that some member is certain it's false?

Advantages: can satisfy both Conditionalization and Anonymity; also compatible with Neutrality... but requires Pr is *uniform* over W.

Disadvantages: Irrelevant Alternatives must be violated (Fact 1); we can't have both Unanimity and Neutrality (Fact 2); It is discontinuous—credence can suddenly drop to 0 from high after a single member drops to 0 from low; If everyone agrees to become more confident (but not certain) in a proposition, the group credence will not change.

IV.2 The Geometric Rule

This next rule is *weird*, but does better. We need two definitions:

- The **geometric mean** of n numbers, $g(x_1, ..., x_n)$, is the nth root of their product; e.g. $g(x_1, x_2, x_3) = \sqrt[3]{(x_1)(x_2)(x_3)}$.
- The **unnormalized group credence** in a world w, $gm(C)(w)$, is the geometric mean of $C(w)$: $gm(C)(w) = \sqrt[n]{c_1(w) \times ... \times c_n(w)}$

(Geometric Rule) For all $w \in W$, $ag(C)(w) = \frac{\sqrt[n]{c_1(w) \times ... \times c_n(w)}}{\Sigma_{w \in W} gm(C)(w)}$

In words: *Take the geometric mean for a world w, and divide it by the sum of geometric means of all the worlds.*

Advantages: Can satisfy Conditionalisation, Neutrality and Anonymity; it is continuous,[13] and is sensitive to agreed upon shifts in credence.

[13] RHB prove that the geometric rule is the *only* rule that can satisfy Conditionalization, Continuity and Neutrality.

Disadvantages: Forces the "pooling" notion of evidence: if just a single person is certain of $\neg A$, then the group credence in A will be 0;[14] Must violate Unanimity or Neutrality; it is "partition-sensitive"—changing what we take to be "worlds" changes the group credence.

[14] Accordingly, the rule is not defined when everybody assigns credence 0 in the same proposition.

References

Hedden, B. (2011). Incoherence without exploitability. *Noûs*, 47(3):482–495.

Lewis, D. (2011). Why conditionalize. In Eagle, A., editor, *Philosophy of Probability: Contemporary Readings*, pages 403–407. Routledge.

Moss, S. (2011). Scoring rules and epistemic compromise. *Mind*, 120(480):1053–1069.