Probabilistic Opinion Pooling

Franz Dietrich
Paris School of Economics / CNRS
www.franzdietrich.net

March 2015
University of Cergy-Pontoise
drawing on work with Christian List (LSE) and work in progress with Marcus Pivato (Cergy-Pontoise)

The problem of group agency:

teating a group as an agent, whith its own beliefs, desires, intensions, actions, plans, ...

Example: the House of Lords

The aggregative approach:

Group attitudes – beliefs, preferences, goals, ... – are not disconnected from people's attitudes, but a function ('aggregation') of them

My focus today: beliefs / opinions

Problem: how merge people's conflicting opinions into group opinions?

Testing your intuition 1

- Ann and Bob wonder whether it will rain tomorrow (they will go hiking).
- Assumption:
  - Ann thinks it will rain with 90% probability;
  - Bob thinks it will rain with 80% probability.
- How probable is rain overall, i.e., from the group's perspective?

Testing your intuition 2

- New start!
- Assumption:
  - Ann thinks rain is 9 times more likely than no rain;
  - Bob thinks rain is 4 times more likely than no rain;
- How much more likely is rain than no rain overall?

Oh!

- The two scenarios are equivalent: Rain is 9 (4) times more likely than no rain *if and only if* rain is 90% (80%) likely.
- **So:** If your first answer was ‘rain is 85% probable’, your second answer should be ‘rain is $\frac{85\%}{15\%} \approx \mathbf{5.666...}$ times more likely than no rain’.

Lesson:

if group beliefs are averages of individual beliefs, then it's unclear what exactly to average

Next test question

- After Ann and Bob tell each other their subjective probabilities of rain (90% and 80%), Ann says:

  "Look, Bob. We both came up independently with high probabilities of rain. This is double confirmation for rain! The overall probability of rain must be at least 95%."

- Who agrees with Ann?

N.B.: exaggerating rather than averaging

- If Ann is right, then the group's overall opinion (probability) should be **more extreme** than each group member's opinion – rather than some **average** or **compromise** between the group members' opinions.

Is Ann right?

This depends, among other things, on Ann's and Bob's sources of information:

- If both drew on similar information (e.g., both heard the same weather forecast), there is no real 'double-confirmation', and Ann is wrong.
- If both drew on independent information, Ann might be right.

Bottom line:

Opinion pooling is non-trivial !!

Three rival approaches

- What should a group's overall probability of a scenario be?
- Three potential answers:
  - an **arithmetic average** of people's probabilities ('linear pooling')
  - a **geometric average** of people's probabilities ('geometric pooling')
  - a **product** of people's probabilities ('multiplicative pooling')[1]

[1] The exact definitions of geometric and multiplicative pooling involve slightly more than taking a geometric average or product. Details soon!

Opinion pooling formalized

The individuals

A group of $n \geq 2$ individuals, labelled $i = 1, ..., n$, who have to assign collective probabilities to some relevant events.

The scenarios (‘possible worlds’)

- Ω : set of possible *worlds*/*states*/*scenarios*/... (Ω is non-empty and finite)
- Examples:
  - $\Omega = \{$rainy, not-rainy$\}$
  - $\Omega = \{$rainy, cloudy, bright$\}$
  - $\Omega = \{$Hollande wins next presidential election, Sarcozy wins, someone else wins$\}$
  - $\Omega = \{0, 1, ..., m\}$, where a ω in Ω represents the number of students coming to my office hours next week
  - ...

Probability functions

- An individual's beliefs/opinions are captured by his probability function.
- A *probability function* is a function P which maps each scenario ω in Ω to a probability $P(\omega) \geq 0$ such that the total probability is $\sum_{\omega \in \Omega} P(\omega) = 1$.

Aggregation (pooling)

- A combination of opinion functions across the n individuals, $(P_1, ..., P_n)$, is called a **profile**.
- A **pooling function** or **aggregation function** is a function F which transforms any profile $(P_1, ..., P_n)$ of individual probability functions into a single collective probability function $P = F(P_1, ..., P_n)$, often denoted $P_{P_1,...,P_n}$.

Linear pooling ("arithmetic averaging")

- The pooling function is **linear** if, for every profile $(P_1, ..., P_n)$, the collective probability of each scenario ω is a weighted arithmetic average

$$P_{P_1,...,P_n}(\omega) = w_1 P_1(\omega) + \cdots + w_m P_n(\omega)$$

  of people's probabilities, for some fixed weights $w_1, ..., w_n \geq 0$ of sum 1.
- **Extreme case:** If $w_i = 1$ for some 'expert' i and $w_j = 0$ for all other individuals j, then we obtain an 'expert rule' given by $P_{P_1,...,P_n} = P_i$.

Geometric pooling ("geometric averaging")

- The pooling function is **geometric** if, for every profile $(P_1, ..., P_n)$, the collective probability of each scenario ω takes the form

$$P_{P_1,...,P_n}(\omega) = c[P_1(\omega)]^{w_1} \cdots [P_n(\omega)]^{w_n}$$

  where $w_1, ..., w_n$ are fixed non-negative weights with sum 1 and c is a scaling factor, given by

$$c = \frac{1}{\sum_{\omega' \in \Omega} [P_1(\omega')]^{w_1} \cdots [P_n(\omega')]^{w_n}}.$$

- **Extreme case:** If $w_i = 1$ for some 'expert' i and $w_j = 0$ for all other individuals j, then we again obtain an 'expert rule' given by $P_{P_1,...,P_n} = P_i$.
- Geometric pooling assumes that $\cap_i \mathrm{supp}(P_i) \neq \varnothing$ (to ensure c is well-defined).

Multiplicative pooling

- The pooling function is **multiplicative** if, for every profile $(P_1, ..., P_n)$, the collective probability of each scenario ω takes the form

$$P_{P_1,...,P_n}(\omega) = cP_1(\omega) \cdots P_n(\omega)$$

where c is a scaling factor, given by

$$c = \frac{1}{\sum_{\omega' \in \Omega} P_1(\omega') \cdots P_n(\omega')}.$$

- Multiplicative pooling assumes that $\cap_i \textsf{supp}(P_i) \neq \varnothing$ (to ensure c is well-defined).

Which of the three pooling methods – linear, geometric, multiplicative – is best?

The *axiomatic method* can help us give answers!

The "indifference preservation" axiom

- The axiom informally: If each individual finds all scenarios equally likely, then so does the group.
- The axiom formally: If each of $P_1, ..., P_n$ is the uniform probability distribution, so is $P_{P_1,...,P_n}$.

$=>$ Satisfied by linear, geometric and multiplicative pooling

The "consensus preservation" axiom

- The axiom informally: If all individuals agree, i.e., have the same beliefs, then these shared beliefs become thee collective beliefs.
- The axiom formally: If $P_1 = \cdots = P_n = P$, then $P_{P_1,\ldots,P_n} = P$.

=> Plausible?

=> Satisfied by linear and geometric pooling, not multiplicative pooling

The axiom of "scenarios-wise pooling"

- The axiom informally: The group's probability of a scenario is determined by people's probabilities of *this* scenario (irrespective of people's probabilities of other scenarios).
  - So the group's probability of "rain" depends only on people's probabilities of rain, not on people's probabilities of "clouds", "sunshine", "hail", "snow", ...
- The axiom formally: The collective probability of a world ω is expressible as a function of people's probabilities of this world, $P_1(\omega), ..., P_n(\omega)$.

$=>$ Plausible?

$=>$ Satisfied by linear pooling only

The axiom of "Bayesianity"

- The axiom informally: If all individuals learn the same event, then group beliefs change by conditionalisation on this event.
- The axiom formally: for any event $E \subseteq \Omega$ (consistent with profile), $P_{P_1^E,\ldots,P_n^E} = P^E_{P_1,\ldots,P_n}$.

$=>$ Satisfied by geometric and multiplicative pooling

The axiom of "external Bayesianity"

- This axiom concerns learning a likelihood function $L : \Omega \to (0, \infty)$, not en event.
- Example: Learning statistical data in Bayesian statistics with multiple statisticians, where Ω is the parameter space.
- For a probability function P and a likelihood function $L : \Omega \to (0, \infty)$, we write P^L for the *posterior probability* function conditional on L, defined as the unique probability function which, as a function of worlds, is proportional to $P \cdot L$.

- The axiom informally: If all individuals learn a likelihood function, then group beliefs change by conditionalisation on it.
- The axiom formally: for any likelihood fn. L, $P_{P_1^L,\ldots,P_n^L} = (P_{P_1,\ldots,P_n})^L$.

$=>$ Satisfied by geometric pooling only

The axiom of "individual-wise Bayesianity"

- Assume only one individual learns the information.
  => idea: people have access to different sources of information ('informational asymmetry')
- The axiom informally: It someone learns a likelihood function, group beliefs change be conditionalisation on it.
- The axiom formally: For any information L and individual i, $P_{P_1,\ldots,P_i^L,\ldots,P_n} = (P_{P_1,\ldots,P_n})^L$.

=> Plausible?

=> Satisfied by multiplicative pooling only

Summary

	linear	geometric	multiplicative
indifference preserving?	X	X	X
consensus preserving?	X	X	
scenario-wise?	X		
Bayesian?		X	X
externally Bayesian?		X	
individual-wise Bayesian?			X

Three theorems

Theorem 1. (Aczél-Wagner 1980, McConway 1981) The linear pooling functions are the only consensus-preserving and scenario-wise pooling functions (assuming Ω contains more than two scenarios).

Theorem 2. (Genest 1984) The geometric pooling functions are unanimity-preserving and externally Bayesian (but there exist other pooling functions with these two properties).

Theorem 3. (Dietrich-List 2014) The multiplicative pooling function is the only indifference-preserving individual-wise Bayesian pooling function.

Our question "which of the three pooling functions is best?" has been reduced to another question: "which axioms are appropriate?"

Which pooling functions and axioms are right?

=> Goal-dependence

- Do we pursue an **epistemic or procedural** goal?
- I.e., should group opinions 'track the truth' or 'track people's opinions'?

Under the epistemic goal, linear pooling looks bad:

- It can't handle information-learning! Neither externally nor individual-wise Bayesian.
- Its central axiomatic property – scenario-wise pooling – lacks an epistemic justification.

Which pooling function and axioms are right?

=> Context-dependence

Suppose the goal is epistemic: we aim for *true* group beliefs!

How to aggregate depends crucially on the informational setting:

Case 1: **symmetric information:** all individuals base their beliefs on the same information

=> Here, geometric pooling and external Bayesianity are plausible.

Case 2: **asymmetric information:** each individual holds private information

=> Here, multiplicative pooling and individual-wise Bayesianity are plausible.

Cases between public and private information

Informational axioms

- "If information is learnt, group beliefs change by conditionalisation on it"
- There are $2 \cdot 3 = 6$ variants of this axiom, depending on
  - what is being learnt, i.e., an event or a likelihood function,
  - who learns the information, i.e., everyone (public info), a single individual (private info) or an arbitrary subgroup.

Conjectures

Bottom line: It depends!

Before pooling opinions, one must know

(i) the goal,

(ii) the informational context.

Thanks!