



Synthese conference 2014
ILLC Amsterdam

Pooling as a Bayesian update: relations with Condorcet and Aumann

★

Jan-Willem Romeijn
University of Groningen
Partly joint work with Olivier Roy (Bayreuth)

Opinion pooling

Raquel and Quassim are both pondering over the proposition A . Raquel's belief is $P_R(A) = r$, Quassim's is $P_Q(A) = q$. Pooling determines that

$$r' = wq + (1 - w)r.$$

The parameter $w \in [0, 1]$ determines how much Raquel moves towards Quassim. It measures the trust of Raquel in Quassim.

Conditionalizing on opinions

A different model for Raquel's accommodating Quassim's opinion employs Bayesian conditionalization on Quassim's opinion:

$$r' = P_R(A|\ulcorner q \urcorner \cap \ulcorner r \urcorner) = P_R(A|\ulcorner r \urcorner) \frac{P_R(\ulcorner q \urcorner|A \cap \ulcorner r \urcorner)}{P_R(\ulcorner q \urcorner|\ulcorner r \urcorner)}.$$

Here Quassim's belief that $P_Q(A) = q$ is denoted by $\ulcorner q \urcorner$, and similarly Raquel's belief is denoted $\ulcorner r \urcorner$. They are categorical events concerning probabilistic opinions.

Bayes, Condorcet and Aumann

We can provide a Bayesian model of opinion pooling. The formal link may be employed to relate pooling to other Bayesian models of epistemic interaction.

- First we focus on the interpretation of the trust parameter w , elaborating its relation to the so-called truth-conduciveness of jurors from Condorcet's theorem.
- Then we relate the representation of pooling as updating to consensus formation and disagreement among peers, in particular the agreement theorem of Aumann.

Contents

1 Pooling as updating	6
2 Truth-conducive voting	9
3 Pooling as voting	11
4 Consensus formation	17
5 The agreement theorem	20
6 Consensus as agreement	26
7 Discussion	29

1 Pooling as updating

Genest and Schervish (1985) establish that we can always find likelihoods $P_R(\lceil q \rceil | A \cap \lceil r \rceil)$ such that, after conditionalizing on $\lceil q \rceil$, Raquel's belief in A equals the result of pooling.

Corrolary of Genest and Schervish (1985)

Let $P_Q(A) = q$, $P_R(A) = P_R(A | \lceil r \rceil) = r$, and let $P'_R(A) = r' = wq + (1 - w)r$ be the result of linear pooling. We choose

$$P_R(\lceil q \rceil | A \cap \lceil r \rceil) = g(q, r) \left(1 - w + \frac{w}{r} q \right),$$

$$P_R(\lceil q \rceil | \neg A \cap \lceil r \rceil) = g(q, r) \left(1 + \frac{r}{1-r} w - \frac{w}{1-r} q \right).$$

Pooling as updating (continued)

Here $g(q, r) = P_R(\lceil q \rceil | \lceil r \rceil)$ is such that

$$\int_0^1 g(q, r) dq = 1,$$

$$\int_0^1 qg(q, r) dq = r.$$

Then the Bayesian update on $\lceil q \rceil$ is identical to the update by linear pooling, $P_R(A | \lceil q \rceil \cap \lceil r \rceil) = r' = P'_R(A)$.

It is intuitive that the trust parameter w shows up as the skewness of the likelihood function.

Raquel's expectations about Quassim

The interpretation of the constraints on $P(\lceil q \rceil \lceil r \rceil)$ is rather natural: Raquel's distribution for Quassim's opinion q is centred on r .

- Raquel might use a peaked Beta-distribution to express that she thinks Quassim will think much like herself.
- She can use a U-shaped Beta-distribution to express the idea that Quassim will be opinionated.

In all of this it is assumed that the events $\lceil q \rceil$ form a partition. We can drop constraints by presuming Quassim might return a blank.

2 Truth-conducive voting

In Condorcet's setting, jurors are asked to vote for or against a proposition A , denoted by V and $\neg V$ respectively. It is standardly assumed that the jurors are competent:

$$c_A = P(V|A) > \frac{1}{2} > P(V|\neg A),$$

$$c_{\neg A} = P(\neg V|\neg A) > \frac{1}{2} > P(V|\neg A).$$

Under this assumption we can derive that for ever larger juries, the majority vote is ever more probable to be correct.

Truth-conducive jurors

We can drop the constraint on the absolute competence of the jurors and assume that jurors have a positive truth-conduciveness Δ :

$$P(V|A) - P(V|\neg A) = c_A + c_{\neg A} - 1 = \Delta > 0,$$

$$P(\neg V|\neg A) - P(\neg V|A) = \Delta > 0.$$

Under this assumption we can still derive a Bayesian version of the Condorcet theorem: votes for A increase the posterior of A , and votes against A decrease it.

3 Pooling as voting

We relate the Condorcet setting to opinion pooling, by making opinion pooling categorical and by using the Bayesian representation of pooling.

- The juror Quassim casts a categorical vote, but Raquel takes him to express a probabilistic opinion.
- The votes for or against A are captured by the events $\lceil q > r \rceil$ and $\lceil q < r \rceil$.
- Raquel accommodates the coarse-grained opinion by means of Bayesian conditioning, just as in voting.
- For this she uses the marginal likelihood of the events $\lceil q > r \rceil$ and $\lceil q < r \rceil$, as determined by the Bayesian equivalent of pooling.

A simplifying assumption

Raquel distinguishes between Quassim offering a degree of belief in A nearby the extremes, but within these two ranges assumes the distribution to be uniform:

$$P_R(q|r) = \begin{cases} l & \text{if } q < \epsilon r, \\ h & \text{if } q > 1 - \epsilon(1 - r), \\ 0 & \text{else.} \end{cases}$$

Solving for the constraints yields:

$$l = \frac{1 - r}{\epsilon r}, \quad h = \frac{r}{\epsilon(1 - r)}.$$

Truth-conducive pooling

Using the afore-mentioned likelihoods and the uniform priors within $\lceil q > r \rceil$ and $\lceil q < r \rceil$, we can derive:

$$P_R(\lceil q > r \rceil | A) - P_R(\lceil q > r \rceil | \neg A) = w \left(1 - \frac{\epsilon}{2} \right),$$

and the same for $\neg A$. Now notice the similarity with the truth-conduciveness of jurors from Condorcet:

$$P(V|A) - P(V|\neg A) = \Delta.$$

Trust as truth-conduciveness

The formulas involving a trust parameter in pooling can be interpreted as expressing a form of truth-conduciveness. For diminishing ϵ we have

$$\Delta = w.$$

The main conclusion is that we can thus express the trust w intuitively and as internal to the model of beliefs.

Modelling choices

Some considerations on the modelling choices needed to relate trust to beliefs:

- The shape of the distribution $P_R(r|q)$ is not crucial. For smaller ϵ it loses import altogether.
- The translation of a categorical vote into a probabilistic opinion can be motivated by a threshold notion of full belief.
- The literature has several other interpretations of w that are less specific but fit well with the current proposal.
- The link between trust and belief may nevertheless come across as somewhat contrived.

Follow-up research

Opinion pooling has a Bayesian reconstruction, which allows us to connect pooling to voting. This may stimulate research at the intersection of pooling with other disciplines.

- The peer disagreement debate has little contact with formal interactive epistemology. The current paper may serve a constructive role in establishing contact, e.g., by illuminating the strategy of splitting-the-difference in disagreement.
- Another application concerns consensus formation. The consensus formation process of DeGroot (1974) and Lehrer and Wagner (1981) can be modelled as a dynamic approach to common knowledge (Genneakoplos and Polemarchakis 1982).

4 Consensus formation

Opinion pooling can be iterated. Writing $P_R^i(A) = r_i$ and $P_Q^i(A) = q_i$, we have on every round $i > 0$:

$$r_{i+1} = (1 - w_R)r_i + w_R q_i.$$

and similarly for Quassim. If the trust parameters w_R and w_Q are fixed and positive, the result is a convergent series of opinion pairs:

$$\langle r_1, q_1 \rangle, \langle r_2, q_2 \rangle, \dots, \langle r_i, q_i \rangle, \dots, \langle p, p \rangle$$

The same things apply to opinion pools with more than two agents.

Aumann's agreement

Aumann uses a Bayesian model for the interacting agents. He proves that the agreement of opinions is automatic if we assume the opinions to be common knowledge. He writes:

“It seems to me that the Harsanyi doctrine is implicit in much of [the literature on opinion pooling]. . . The result of this paper may be considered a theoretical foundation for the reconciliation of subjective probabilities [i.e., by means of pooling].”

Surprisingly, there is no account of how the agreement theorem relates to iterated opinion pooling.

Agreement and consensus

The remainder of this paper provides a reconstruction of the approach to consensus based on Aumann's result. More precisely:

- The consensus formation can be represented as a dynamic approach to common knowledge.
- The Bayesian rendering of opinion pooling determines the requisite constraints on the common prior.

This rationalizes consensus formation via pooling, thereby revealing its conditions for applicability. Moreover, it offers a new perspective on pooling as information sharing through higher-order knowledge.

5 The agreement theorem

In Aumann's (1976) theorem, Raquel and Quassim are in the following epistemic situation:

- They share a space Ω of possible worlds and an initial probability assignment P over it.
- Both have their own information partition \mathcal{R} and \mathcal{Q} , with elements R_i and Q_j . These partitions are common knowledge.
- Each of them has private information, in the form of one element from their partition, R_0 and Q_0 .

Common knowledge

A central notion of the theorem is that of common knowledge, which relies on a specific conception of knowing:

Knowledge of knowledge of. . .

Raquel knows the proposition X iff R_0 is included in X . She knows that Quassim knows X iff all Q_j that intersect with R_0 , i.e., that can be reached from R_0 , are included in X . And so on.

Common knowledge of the posteriors r and q is associated with a particular set C_{rq} that is included in both $\lceil r \rceil$ and $\lceil q \rceil$.

A static result

Using this notion of common knowledge, Aumann proves the following.

Theorem

If two people have the same priors, and their posteriors for an event A are common knowledge, then these posteriors are equal.

We can divide the set C_{rq} using the partitions \mathcal{R} and \mathcal{Q} . For all R_i and Q_j overlapping with C_{rq} :

$$P(A|R_i \cap C_{rq}) = r \quad \text{and} \quad P(A|Q_j \cap C_{rq}) = q.$$

It does not matter in what direction we marginalize the probability for A , hence $r = q$.

Reaching agreement

Geanakoplos and Polemarchakis (1982) provide an account of how agreement may arise from an exchange of opinions between Raquel and Quassim.

- At the outset both receive private information, R_0 and Q_0 , from their own information partition.
- At each round i they exchange their posteriors r_i and q_i .
- With this new information they exclude members from the information partition of the other agent.
- And they update their own opinions accordingly.

Sharing information

In the first round Raquel conditions on Quassim's opinion $\lceil q_1 \rceil$ and obtains a new opinion about A :

$$P_R^2(A) = P_R^1(A|\lceil q_1 \rceil) = P(A|\lceil q_1 \rceil \cap R_0).$$

The set $\lceil q_1 \rceil$ comprises the elements from Quassim's information partition that are consistent with the probability assignment $P_Q^1(A) = q_1$:

$$\lceil q_1 \rceil = \{Q_j : P(A|Q_j) = q_1\}.$$

Quassim also deletes everything outside $\lceil q_1 \rceil$ from the sets of Q_j that Raquel considers possible. Raquel and Quassim similarly update on $\lceil r_1 \rceil$.

Higher-order knowledge

The process is then iterated on the resulting smaller space. Notice that Raquel and Quassim update on events that sit ever higher up in a hierarchy of knowledge,

$$\lceil q_j \rceil = \bigcup \{Q_k : P(A|Q_k \cap \lceil r_{j-1} \rceil) = q_j\} \cap \lceil q_{j-1} \rceil \cap \lceil r_{j-1} \rceil,$$

and similar for r_i . A more explicit representation of this can be given in so-called Harsanyi type space but for present purposes we may collapse type space onto the information partitions.

6 Consensus as agreement

Iterated pooling and reaching agreement manifest in similar ways:

$$\langle r_1, q_1 \rangle, \langle r_2, q_2 \rangle, \dots, \langle r_i, q_i \rangle, \dots, \langle p, p \rangle$$

We can determine the prior over $\mathcal{R} \vee \mathcal{Q}$ such that approaching agreement fits the iterated pooling process. The constraints are

$$P(A|Q_k \cap \lceil r_{j-1} \rceil) = q_j$$

for all j and for all k such that $Q_k \cap \lceil q_{j-1} \rceil$ non-empty. A similar constraint set must hold for r_i .

Main result

Using Genest and Schervish (1985), this set of constraints can be imposed coherently onto the partitioned space Ω .

- The constraints leave some space for variation, especially in the distributions that agents choose for the opinions of other agents.
- Alternative pooling operations will lead to different constraints, e.g., on what opinions we expect others to have.
- It is not part of the model that agents can derive the values of trust parameters from the revealed opinions.
- The assignments within $\lceil q_j \rceil^C$ are determined by the pooling that would happen if Quassim had revealed something other than q_j .

Multiple agents

The result generalizes to any number of agents, whether they learn the other opinions sequentially or all at the same time. Quassim and Simone, say, can be viewed as constituting a single agent:

$$r_{i+1} = (1 - (w_Q + w_S)) r_i + (w_Q + w_S) \left(\frac{w_Q q_i + w_S s_i}{w_Q + w_S} \right).$$

However the model for multiple agents is organized, the likelihoods will encode specific assumptions about dependencies among the agents:

$$P(\lceil q_i \rceil \cap \lceil s_i \rceil | A \cap \lceil r_i \rceil) \neq P(\lceil q_i \rceil | A \cap \lceil r_i \rceil) \times P(\lceil s_i \rceil | A \cap \lceil r_i \rceil).$$

7 Discussion

The result constitutes a bridge between two models of epistemic interaction. It seems natural to transport things over the bridge.

- Can we view pooling as a kind of information sharing? There seem to be conceptual differences between consensus and agreement.
- How can we interpret the common priors? They specify the conditions under which the epistemic shortcut of pooling is warranted.
- In addition, we may look for a taxonomy of consensus formations and consensus failures in terms of the common prior, or lack thereof.

A new angle on pooling

The standard view is that in the context of agreement, the exchange of opinions constitutes implicit information sharing, whereas in pooling the exchange amounts to a series of concessions.

- Consider the evidence presented by the opinions: after the initial round, the information is in opinion *changes*, i.e., in responses of the agents to each other.
- In the agreement context, the opinions refer to ever higher orders of knowledge. We may reinterpret the exchange in the context of pooling along similar lines.

The model of consensus can be viewed afresh, in terms of the conception of opinion as evidence stemming from the agreement context.

Pooling as shortcut

Following Genest and Schervish (1985), the agreement version of consensus by iterated pooling provides insights into the implicit assumptions of pooling.

- Pooling operations ignore many aspects of the distributions in the Bayesian model. We can motivate a consensus formation process by determining what aspects are relevant.
- Pooling entails specific dependencies among the opinions of agents. The Bayesian model offers insights into these dependencies, and this might help to motivate them.

Thank you

The slides of this talk will be available at <http://www.philos.rug.nl/romeyn>.
Papers are available upon request. For comments and questions, email j.w.romeijn@rug.nl.