

Voting, Deliberation and Truth

Stephan Hartmann and Soroush Rafiee-Rad

Tilburg Center for Logic and Philosophy of Science

Episteme@CMU Conference

Pittsburgh, June 2011

Motivation

- A group has to make **yes-no decision**.
- Example: A jury in court has to decide whether the defendant is guilty or innocent.
- There are two ways to make such a collective decision:
 - 1 **Voting**: Cast a vote, and count the number of yes-votes. If it is greater than a certain threshold (e.g. 50%), then the group decides 'yes'.
 - 2 **Deliberation**: The jury members deliberate the issue in question, and will eventually arrive at a consensual decision.
- **Question**: Which of these two procedures is better?
- To address this question, we have to **model** the process of deliberation and specify a **criterion of comparison**.

1. Voting

- Everybody casts a vote, and the proposal is accepted if the number of votes is greater than a certain threshold, e.g. 50 %.
- This procedure is comparatively easy to conduct and feasible for large groups.
- However, voting leads to a **compromise** as not everybody will endorse the group decision and be happy with it.
- Note, though, that observing that the majority of the members of my group does not hold my view might make me change my mind.

2. Deliberation

- In a deliberation, the jury members (try to) convince each other, they exchange arguments and change their views.
- If all goes well, they will eventually arrive at a **consensual decision** which makes everybody happy as everybody endorses it.
- While voting is easier to implement, a deliberation is often considered to be more satisfactory and preferable on procedural grounds.
- Deliberation is only feasible in small groups.
- It is much harder to model than the voting procedure.

- Note: In our example, it is a **matter of fact** as to whether the defendant is guilty or not.
- The jury should make the right decision, i.e. it should decide 'guilty' if the defendant is guilty, and 'not guilty' if the defendant is not guilty.
- This suggests to provide an **epistemic analysis** of the two procedures and to explore which procedure does best in terms of truth-tracking.

- Error Minimization:** We want the jury to (i) minimize the false positives (i.e. $P(\text{'guilty'}|\text{not} - \text{guilty})$) and (ii) to minimize the false negatives (i.e. $P(\text{'not} - \text{guilty'}|\text{guilty})$). These two probabilities characterize the **reliability** of a jury member.
 - Truth Tracking:** We want the jury (or the judge, who considers the vote of the jury) to apply an aggregation procedure that maximizes the probability of making the right decision. That is, we want a decision procedure that does best epistemically.
- The **goal** of this talk is to provide a comparative analysis of voting and deliberation.
 - To do so, we **develop a Bayesian model of deliberation**.

- Voting**
 - The Condorcet Jury Theorem
- Deliberation**
 - A Bayesian Model of Deliberation
 - Deliberation and Truth-Tracking
- Voting vs. Deliberation**
 - Which Procedure Is the Better Truth-Tracker?
 - Homogeneous and Inhomogeneous Groups
- Outlook**

Let there be a group of n voters, which has to make a yes-no decision on a proposition H . We assume:

- Independence:** Given the truth or falsity of H , the verdict of one voter does not depend on the verdict of any other voter.
- Reliability:** Each voter has a certain reliability $r := P(\text{Vote}_Y|Y) = P(\text{Vote}_N|N) > .5$ to make the right decision. That is, we assume that the rate of false positives equal the rate of false negatives.

Then the probability that the majority makes the right decision (i) increases monotonically and (ii) goes to 1 as $n \rightarrow \infty$. (It is actually enough that the average reliability of all voters is greater than .5.) Hence, voting is (given the above conditions) truth-conducive.

The Bayesian Model in a Nutshell

- 1 A group of n members (labeled A_1, \dots, A_n) deliberates, in various rounds, the truth or falsity of a proposition H .
- 2 Everybody assign a prior probability $P_i^{(0)}(H)$ and votes (if sufficiently certain) according to this probability assignment. The votes are announced.
- 3 Everybody updates her probability assignment, taking the votes of the others into account.
- 4 Iterate this.
- 5 Cast a vote, if no consensus emerges after a certain number of rounds.

Question: How to weigh the verdicts of the other group members?

Details I: Voting

- We introduce the propositional variable V_i for the vote of group member A_i . The value V_i means that A_i votes for H , and the value $\neg V_i$ means that A_i votes for $\neg H$.
- To cast a vote means to map each A_i 's prior into 'yes' or 'no'. To do so, we apply a chance mechanism and choose a random number $t \in (0, 1)$ from a uniform distribution.

$$\begin{aligned} \text{Vote}_i^{(0)} &= V_i, \text{ if } t \leq P_i^{(0)}(H) \\ \text{Vote}_i^{(0)} &= \neg V_i, \text{ otherwise.} \end{aligned}$$

- Furthermore, we introduce the variables $p_i^{(0)}$ and set $p_i^{(0)} = 1$ if $\text{Vote}_i^{(0)} = V_i$ and $p_i^{(0)} = -1$ if $\text{Vote}_i^{(0)} = \neg V_i$.

Details II: Two Reliabilities

Each group member A_i is characterized by two reliabilities:

- 1 A first order reliability (r_i): The probability to make the correct judgment regarding the truth of H .
- 2 A second order reliability ($c_i^{(0)}$): The probability to assess the first order reliability of another group member.

Details III: Assessing Others

- Each group member A_i assigns a reliability value $r_{ij}^{(0)} := P_i^{(0)}(V_j|H) = P_i^{(0)}(\neg V_j|\neg H)$ to group member $A_j \neq A_i$.
- The value of $r_{ij}^{(0)}$ depends on $c_i^{(0)}$ and r_j . It follows from a (suitably transformed) β -distribution over the interval $[\max(0, r_j + c_i^{(0)} - 1), \min(1, r_j - c_i^{(0)} + 1)]$ around r_j .
- Clearly, the higher the second order reliability, the better the estimate of someone else's first order reliability.
- We use these (estimated) reliabilities to calculate the ...

Likelihood ratios.

$$x_{ij}^{(0)} := \frac{P_i^{(0)}(V_j|\neg H)}{P_i^{(0)}(V_j|H)} = \frac{1 - r_{ij}^{(0)}}{r_{ij}^{(0)}}$$

Details IV: Updating

- The prior probability is updated taking the votes of the other group members into account. To do so, we assume conditional probabilistic independence (as in the case of the Condorcet Jury Theorem), i.e. $Vote_i^{(0)} \perp\!\!\!\perp Vote_j^{(0)} | H$.
- With this assumption, we calculate, using Bayes' Theorem:

The posteriors

$$P_i^{(1)}(H) := P_i^{(0)}(H | Vote_1^{(0)}, \dots, Vote_{i-1}^{(0)}, Vote_{i+1}^{(0)}, \dots, Vote_n^{(0)})$$

$$= \frac{P_i^{(0)}(H)}{P_i^{(0)}(H) + (1 - P_i^{(0)}(H)) \prod_{k=1, k \neq i}^n (x_{ik}^{(0)})^{P_k^{(0)}}$$

Details V: Iteration and Update of the $c_i^{(0)}$'s

- In the second round of deliberation, each group member will use $P_i^{(1)}(H)$ as her new prior.
- In addition, we assume that the second order reliabilities increase linearly during the course of deliberation until a maximum value of $C \leq 1$ is reached, i.e.

$$c_i^{(k)} = c_i^{(0)} + \frac{k \cdot (C - c_i^{(0)})}{M}$$

- Here M is the maximum number of deliberation rounds.
- If there is no consensus after M rounds, then majority voting is applied.

Does Deliberation Track the Truth?

- We explore our model using computer simulations.
- We first ask whether deliberation is a truth-conducive procedure. To do so, we calculate the probability that the procedure gives the correct answer as a function of the group size.
- More specifically, we investigate two types of groups:
 - 1 Homogeneous groups: all the group members have the same first order reliability.
 - 2 Inhomogeneous groups: not all group members have the same first order reliability.

Deliberation Is Truth-Conducive: Homogeneous Groups

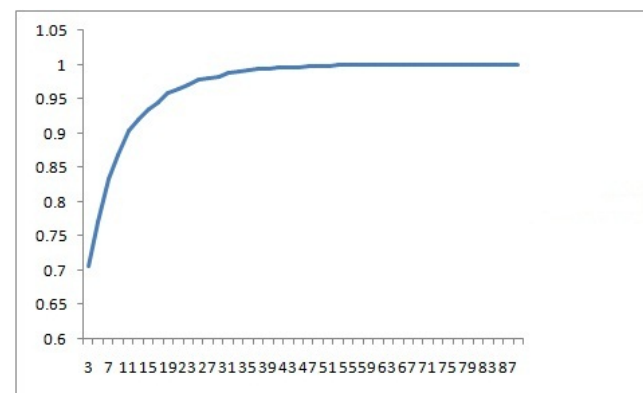


Figure: All the group members have a reliability of .7.




Figure: One third of the group members have a reliability of .8, two thirds have a reliability of .6.




Figure: The probability that deliberation tracks the truth as a function of $c_i^{(o)}$ for a group of size 15. Five group members have a reliability of .6 and ten have a reliability of .75.

- Having (numerically!) established that deliberation is also truth-conducive, we now compare which of our two procedures does better in terms of truth tracking.
- To do so, we calculate

The difference Δ

$$\Delta = P_D - P_V,$$

i.e. the probability that the deliberation process tracks the truth minus the probability that the voting procedure tracks the truth, as a function of the group size.

- To compare the two procedures, we make sure that various parameters (group size, reliabilities) are the same for both procedures.




Figure: Δ as a function of the group size for homogeneous groups and for different reliabilities.

- For homogeneous groups, voting always does better than deliberation.
- This result is hardly surprising as the weighted average, of which the majority rule for voting is a special case (i.e. all voters get the same weight), has been shown to be epistemically optimal. See Nitzan and Paroush (1982) and Gradstein and Nitzan (1986).
- That is, if one knows that the group is homogeneous, or if one wants to consider the group to be homogeneous (for political or whatever reasons), then majority voting does best.

- Note that the theoretical result that the weighted average performs best epistemically also applies to inhomogeneous groups, and one obtains the epistemically optimal result by assigning weights which are proportional to the corresponding reliabilities.
- However, even if one would know these reliabilities, it would not be possible for political reasons to use them in the aggregation procedure (e.g. to make them known).
- And so straight voting or something like our deliberation procedure will typically have to be used.

- For inhomogeneous groups, we distinguish two cases:
- 1 Almost homogeneous groups: The reliabilities of all group members are close-by.
 - 2 Manifestly inhomogeneous groups: The reliabilities of all group members differ considerably.




Figure: Δ as a function of the group size for almost homogeneous groups and for different reliabilities.




Figure: Δ as a function of the group size for almost homogeneous groups and for different reliabilities.

We conclude that for almost homogeneous groups the situation is similar to the situation for homogeneous groups, i.e. voting has a higher chance of arriving at the truth than deliberation.




Figure: The larger portion of the group has a higher reliability.




Figure: The larger portion of the group has a higher reliability.




Figure: The larger portion of the group has a lower reliability.




Figure: The larger portion of the group has a lower reliability.

We conclude that for manifestly inhomogeneous groups, deliberation has a higher chance of arriving at the truth than voting, in particular for groups of small and medium size.

This is even the case if only one member of the group has a considerably different reliability, as the following figure shows.




Figure: One group member is highly reliable (reliability = .9). The rest has a reliability of .6.




Figure: Δ as a function of the second order reliability $c_i^{(0)}$. Two thirds of the group have a reliability of .6. The rest has a reliability of .75.




Figure: Δ as a function of the second order reliability $c_i^{(0)}$ for a group of size 21 and different reliability distributions.

- 1 Deliberation gives better results for high second order reliabilities.
- 2 Voting gives better results for low second order reliabilities.
- 3 The turning point depends on the group size: The larger the group, the lower the threshold for the second order reliabilities.

- 1 Voting and deliberation are, given certain conditions, truth-conducive.
- 2 It depends on the context, which procedure is the better truth-tracker.
- 3 For homogeneous and almost homogeneous groups, voting has a higher chance of arriving at the truth.
- 4 For inhomogeneous groups with high second order reliability, deliberation works better.

Open questions: Study network effects, include arguments in the model, extend the model to judgment aggregation, and relate the model to empirical studies.

. . . thanks for your attention!