Deliberation and the Wisdom of Crowds

Franz Dietrich (presenter)
Paris School of Economics & CNRS

Kai Spiekermann
LSE

"Social Choice under Risk and Uncertainty"
Warwick Business School
9 May 2022

The paper is downloadable from www.franzdietrich.net

Deliberation vs aggregation

Collective decisions are often reached in two steps:

(1) A deliberation process shapes group members' opinions.

(2) A voting rule merges opinions into a decision.

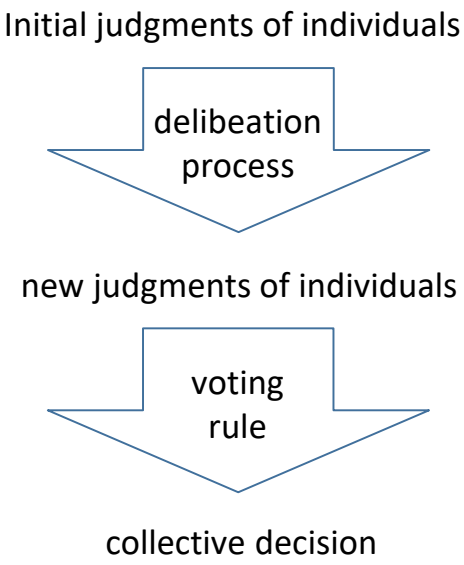

This gives us two instruments for improving group decisions:

- improve the deliberation process, to improve voter opinions
- improve voting rule, to improve the use of voter opinions

The epistemic perspective

But what is a 'good' opinion, decision, or voting rule?

- Epistemic vs procedural criteria for evaluating decisions and opinions
- Epistemic vs procedural criteria for evaluating deliberation/voting processes
- Judgment voting vs preference voting

Our perspective: epistemic all the way

Deliberation & voting in context

An interdisciplinary literature has taken different perspectives on (the epistemics of) deliberation and voting:

- **Social-choice theoretic**
  - jury theorems; powerful, but focused on voting rather than deliberation
- **Game-theoretic**
  - provides micro-foundations; but assumes a rather special (purely strategic) concept of deliberation and voting
- **Normative-democratic**
  - philosophical, not formal

–> see paper for discussion and citations

Our question: does deliberation improve voting outcomes?

Not obvious, because deliberation has rival effects:

- a positive effect of raising voter competence (as one may hope)
- a negative effect of raising voter interdependence (as one may fear)

Which effects prevails?

Our modelling approach

- Our approach is formal, but macroscopic rather than game-theoretic, focusing on the structure of information flow rather than the micro-foundations of behaviour.
- Normally, formal analyses of deliberation are game-theoretic; but see two semi-game-theoretic models of deliberation, namely Chung and Duggan's (2020) model of myopic discussion, constructive discussion and debate, and Ding and Pivato's (2021) model of deliberation as information disclosure.

Why not a game-theoretic model?

Two motivations:

1. **Achieving parsimony**
2. **Permitting several psychological interpretations and hypotheses.** Limited sharing and absorbing of sources could represent either conscious choices or unsuccessful attempts or hard inabilities; and deliberators could be instrumentally or intrinsically motivated, have stable or variable preferences, reason strategically or not, be fully rational or use simple heuristics, anticipate all contingencies or fail to imagine them before they occur during deliberation.[1] NB: Game-theoretic models need to commit on all these issues, often in stylised and specific ways[2]

[1] Relatedly, deliberators could acquire only new information or also refined awareness and concepts to perceive or interpret the situation.

[2] For instance, the game tree describing all contingencies is known to all players. This implies that only information can grow, not awareness.

Goals in a nutshell

We

- consider a group decision problem between two options,
- model deliberation as sharing and absorbing evidences,
- prove (the first ?) jury theorems that address the epistemic benefits of pre-voting deliberation,
- identify 3 voting failures, and analyse whether deliberation can reduce them.

Plan

1. Three voting failures
2. Opinion formation based on sources
3. Deliberation as sharing and absorbing
4. A pre-deliberation and a post-deliberation jury theorem
5. Measuring Failures 1 and 2
6. A typology of beneficial and harmful deliberation
7. Generalised opinion structures and deliberation processes
8. A theoretic analysis of Failure 3

Part 1

Three voting failures

The model, informally

- Voters 1, ..., n.
- Two options (e.g., convict or acquit the defendant)
- One option is 'correct', the other 'incorrect'.
- Each voter accesses some *evidences*, with more or less overlap across voters.
  - NB: Our notion of 'evidence' is broad: empirical facts, arguments, perspectives, or other reasons.
- Each voter forms an opinion about what option is correct, based on *her* evidences.
- The opinions are aggregated though majority rule, leading to a group opinion.

Failure 1: overcounting widespread evidence

- **The Failure:** Evidence held by more voters has exaggerated influence, by affecting more votes.
- **Example:** Evidence 1 affects two votes while evidences 2 and 3 each affect only one vote.

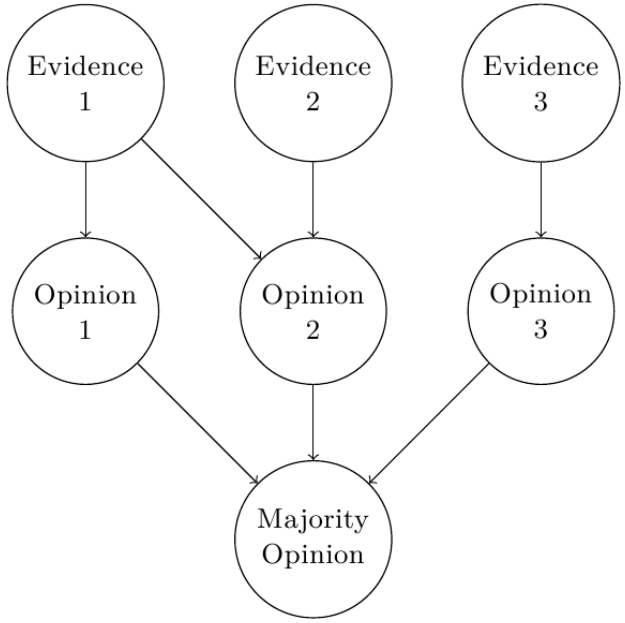

- **Hypothesis**: Deliberation reduces Failure 1 by increasing the spread of previously private or almost private evidence.

Failure 2: neglecting evidential inequality

- **The Failure:** Voters have the same weight, despite their unequally strong total evidence.
- **Example:** Voter 2 has stronger total evidence.[3]

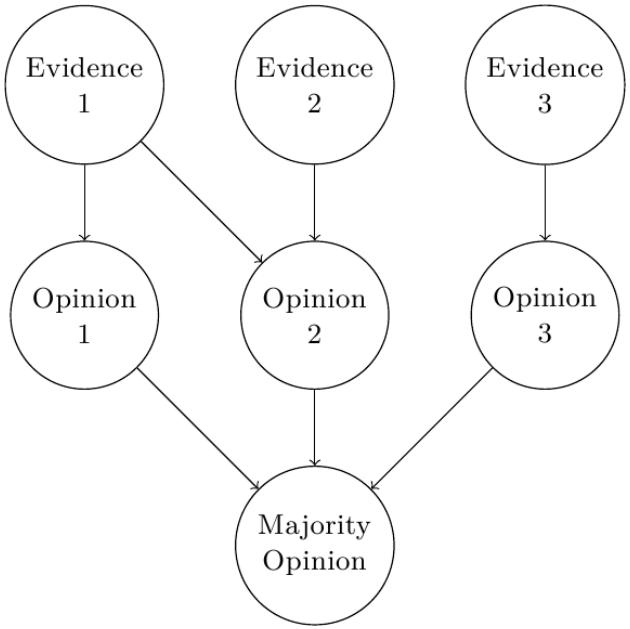

- **Hypothesis**: Deliberation reduces Failure 2 by letting voters with initially weak evidence accumulate evidence.

[3]Assuming her two evidences are stronger in total than evidence 1 and than evidence 3.

Failure 3: neglecting evidential complementarity

- **The Failure:** Information obtainable after combining different evidences dispersed across voters is undercounted, because few or no voters access all these evidences simultaneously.
- **Example:** No voter has both evidences 1 and 3, which are complementary (they might be two arguments that are uninformative in isolation but conclusive in combination).

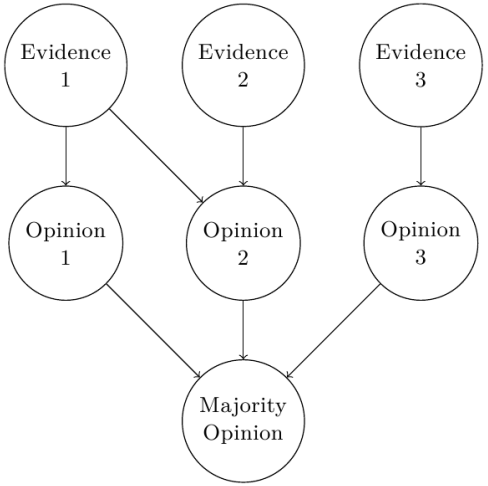

- **Hypothesis**: Deliberation reduces Failure 3 by letting voters collect evidences from others, and then recognize and use evidential complementarities.

Part 2:

Opinion formation based on sources

Persons, options, state

- Persons $1, ..., n$ ($n \geq 2$). Let $N = \{1, ..., n\}$.
- Options 1 and -1.
- One option is (objectively or intersubjectively) correct or better. It is called the *state (of the world)*, represent it by a random variable $\mathbf{x}$ taking the value 1 or -1.
  - Ex: judges decide between 'convict' or 'acquit'; the state is determined by whether the defendant is guilty or innocent.

Notation

- Random variables in bold, their values in non-bold.
- All random variables and events are defined relative to a fixed probability space, with probability measure denoted Pr.

Evidences and their sources

- Opinions are based on *evidences* from *sources*. (The generalised model presented later also allows non-evidential influences.)
- S : a finite non-empty set of *sources (of evidence)*
- $\mathbf{e}_s$: evidence from source s ($\in S$), formally a real-valued random variables.
  - Positive evidence supports option 1, negative option -1, to a strength represented by the absolute value of the evidence.
- Ex: the source s is a witness report, the evidence $\mathbf{e}_s$ measures the evidential support from this source.

Source access and opinions

- Each person i accesses some set of sources, her *source set*, represented by a random variable $\mathbf{S}_i$.
  - Ex: one juror's source set contains some witness report and some legal text interpreting the law, another juror's source set contains the defendant's facial expression etc.
- The *opinion of person* i is the option supported by i's total evidence:

$$\mathbf{o}_i = \begin{cases} 1 & \text{if } \sum_{s \in \mathbf{S}_i} \mathbf{e}_s > 0 \\ -1 & \text{if } \sum_{s \in \mathbf{S}_i} \mathbf{e}_s < 0 \\ 0 & \text{if } \sum_{s \in \mathbf{S}_i} \mathbf{e}_s = 0. \end{cases}$$

- The *competence* of i is the probability of a correct opinion $p_i = Pr(\mathbf{o}_i = \mathbf{x})$.

The collective opinion

- The *majority opinion* is

$$\mathbf{o}_{maj} = \begin{cases} 1 & \text{if } |\{i : \mathbf{o}_i = 1\}| > |\{i : \mathbf{o}_i = -1\}| \ (\Leftrightarrow \sum_i \mathbf{o}_i > 0) \\ -1 & \text{if } |\{i : \mathbf{o}_i = 1\}| < |\{i : \mathbf{o}_i = -1\}| \ (\Leftrightarrow \sum_i \mathbf{o}_i < 0) \\ 0 & \text{if } |\{i : \mathbf{o}_i = 1\}| = |\{i : \mathbf{o}_i = -1\}| \ (\Leftrightarrow \sum_i \mathbf{o}_i = 0). \end{cases}$$

- The *majority competence* is the probability of a correct majority opinion $p_{maj} = P(\mathbf{o}_{maj} = \mathbf{x})$.

Diversity as difference in sources across persons

Terminology

- The *source profile* is the combination of source sets across people $(\mathbf{S}_i)_{i \in N}$, in short $(\mathbf{S}_i)$.
- Person i's *evidence bundle* is the family of evidences from her sources $(\mathbf{e}_s)_{s \in \mathbf{S}_i}$; it is doubly random, through her source set $\mathbf{S}_i$ and the evidences $\mathbf{e}_s$ from her sources s.
- The *evidence profile* is the combination of evidence bundles across people $((\mathbf{e}_s)_{s \in \mathbf{S}_i})_{i \in N}$, in short $((\mathbf{e}_s)_{s \in \mathbf{S}_i})$.

Three assumptions

For now, we make three (strong) assumptions:

- *Equiprobable States*: the state $\mathbf{x}$ takes both values 1 and -1 with probability $\frac{1}{2}$.
- *Simple Gaussian Evidences:* Given any state $x \in \{\pm 1\}$, the evidences $\mathbf{e}_s$ ($s \in S$) have independent Gaussian distributions with mean x and some variance σ^2 that is the same across states x and sources s.[4]
- *Independent Sources*: The source-access events[5] are independent across persons and sources, and jointly independent of the state and the evidences.

[4]Under generalised Gaussian assumptions, the evidences are possibly dependent (given the state), with means and/or variances that can vary across states and/or sources.

[5]i.e., the events 'person i has source s' ('$s \in \mathbf{S}_i$'), for different $i \in N$ and $s \in S$.

Access probabilities

- The 'access probabilities' $p_{s \to i} = Pr(s \in \mathbf{S}_i)$ ($s \in S$, $i \in N$) fully determine the distribution of the source profile $(\mathbf{S}_i)$.
- How? By Independent Sources, the probability that person i has source set S_i is

$$Pr(S_i) = \left(\prod_{s \in S_i} p_{s \to i} \right) \left(\prod_{s \in S \setminus S_i} \overline{p_{s \to i}} \right),$$

  and the probability of an entire source profile (S_i) is the product $\prod_i Pr(S_i)$.

Summary

Our formal primitive is a *simple opinion structure*, by which we mean a triple $(\mathbf{x}, (\mathbf{e}_s)_{s\in S}, (\mathbf{S}_i)_{i\in N})$, in short $(\mathbf{x}, (\mathbf{e}_s), (\mathbf{S}_i))$, containing:

(1) a random variable $\mathbf{x}$, the *state* or *correct option*, taking the value 1 or -1 with equal probability;

(2) a family $(\mathbf{e}_s)$, indexed by some set S of *sources* (non-empty and finite), consisting of real-valued random variables, the *evidences* from these sources, which have state-conditionally independent Gaussian distributions with mean the state and with some fixed variance $\sigma^2 > 0$;

(3) a family $(\mathbf{S}_i)$, indexed by some set $N = \{1, ..., n\}$ of *persons* ($2 \leq n < \infty$), consisting of random subsets of S, the *source sets* of these persons, with distributions determined by access probabilities $(p_{s\to i})_{s\in S, i\in N}$ via (**??**), independently across persons and independently of the state and the evidences.

Rationality?

- Is our simple opinion model ad hoc from a rationality perspective?
  - A natural worry, as we presuppose a seemingly native rationale for forming opinions: you add up your evidences, and compare the sum with zero.
- In fact, opinion formation *is* rational, in a perfectly classical sense. Why?

Rationality!

- Standard rationality requires evaluation by expected utility. Given our epistemic setting, let 'utility' mean 'correctness level', defined as 1 if the opinion is correct, 0 if it is incorrect, and $\frac{1}{2}$ if it is neutral, i.e., 0.
- The opinion $\mathbf{o}_i$ is *classically rational* if its expected correctness level is at least that of any possible opinion, i.e., of any random variable $\mathbf{o}'$ that generates 1, -1 or 0 as some function of i's evidence bundle $(\mathbf{e}_i)_{i \in \mathbf{S}_i}$.

Theorem 1 *Under any simple opinion structure* $(\mathbf{x}, (\mathbf{e}_s), (\mathbf{S}_i))$*, the opinion* $\mathbf{o}_i$ *of every person* i *is classically rational.*

An idea of the proof

Theorem 2 *(restated) Under any simple opinion structure* $(\mathbf{x}, (\mathbf{e}_s), (\mathbf{S}_i))$*, the opinion* $\mathbf{o}_i$ *of every person* i *is classically rational.*

Proof idea: All three 'simplifying assumptions' crucially enter this result:

(1) Equiprobable States rationalises that the 'cut-off threshold' is 0.

(2) Simple Gaussian Evidences: rationalises adding (rather than otherwise aggregating) one's evidences. The Gaussian distribution is a miracle in many ways!

(3) Independent Sources ensures that a voter i only learns something from her evidences, not from the fact *that* she knows them (formally, only the evidences $\mathbf{e}_s$ carry information, not the source set $\mathbf{S}_i$ itself). ■

Part 3

Deliberation as sharing and absorbing

Illustration

- In a deliberation process, persons transmit some of their sources to some other persons.
- Example:

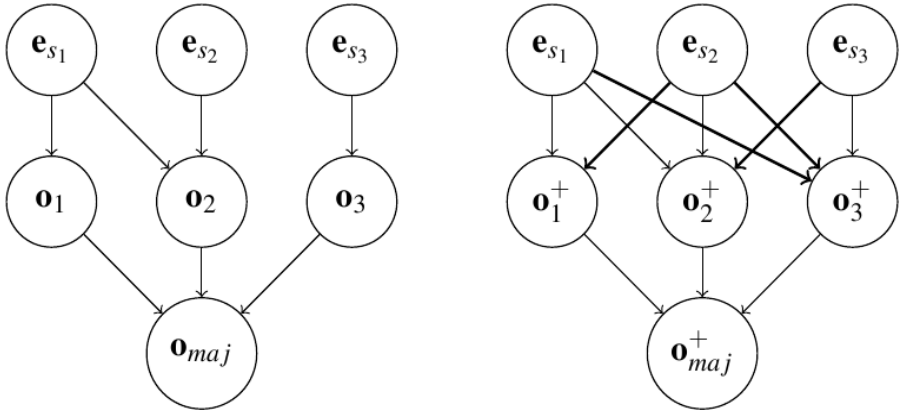

(a) Pre-Deliberation (b) Post-Deliberation

Bold arrows indicate newly absorbed sources.

Share-absorb processes

- For now, we focus on a special type of deliberation process: *share-absorb processes*.
- Such a process is given by 'sharing probabilities' $p_{s,i\rightarrow} \in [0,1]$ and 'absorbing probabilities' $p_{s,i\leftarrow} \in [0,1]$ for all source s and person i.
- Starting from an initial source sets $(\mathbf{S}_i)_{i\in N}$,

(step 1) each person i shares each of her initial sources $s \in \mathbf{S}_i$ with an independent probability of $p_{s,i\rightarrow}$, and then

(step 2) for each source s shared by at least someone, each person i with $s \notin \mathbf{S}_i$ absorbs this source with an independent probability of $p_{s,i\leftarrow}$.

The result of a share-absorb process

- The process transforms the initial source profile $(\mathbf{S}_i)$ into new one $(\mathbf{S}_i^+)$, in which $\mathbf{S}_i^+ = \mathbf{S}_i \cup \{s \in S : i \text{ was absorbed by } i\}$.
- This yields a new opinion structure $(\mathbf{x}, (\mathbf{e}_s), (\mathbf{S}_i^+))$, and thus new personal opinions $\mathbf{o}_i^+$ and competence levels p_i^+, and a new group opinion $\mathbf{o}_{maj}^+$ and group competence level p_{maj}^+.

The hope

- The hope is that deliberation is successful:
  - higher group competence: $p^{+}_{maj} > p_{maj}$
  - smaller Failures 1, 2, and 3
- Mechanism at work: sources have become more accessible.

Variable sharing propensity

- Reasons why the sharing probability $p_{s,i\rightarrow}$ may be source-dependent:
  - easy vs. hard to communicate sources;
  - consciously vs. subconsciously held sources.
  - norms
  - etc.
- Reason for person-dependent sharing probabilities:
  - better or worse communicators
  - variable motivation
  - variable attention
  - etc.

Variable absorbing propensity

- Reasons for source-dependence of absorbing probability $p_{s,i\leftarrow}$:
  - easy vs. hard to understand
  - etc.
- Reasons for person-dependence:
  - open-minded vs. stubborn
  - able vs. unable to catch new arguments
  - etc.

Part 4

A pre-deliberation and a post-deliberation jury theorem

Preview

- The wisdom of crowds is often defended through jury theorems, typically with the optimistic conclusion that larger groups perform better and better, and ultimately become infallible.
- In our framework, a less optimistic jury theorem holds: the wisdom of crowds is objectively bounded, even in asymptotically large groups.
- But post-deliberation things look better.

The objective upper bound

- The best the group can hope for is the *ideal competence*, based on the total evidence.
- The *ideal opinion* is based on all sources:

$$\mathbf{o}_{IDEAL} = \begin{cases} 1 & \text{if } \sum_{s \in S} \mathbf{e}_s > 0 \\ -1 & \text{if } \sum_{s \in S} \mathbf{e}_s < 0 \\ 0 & \text{if } \sum_{s \in S} \mathbf{e}_s = 0. \end{cases}$$

- The *ideal competence* is

$$p_{IDEAL} = Pr(\mathbf{o}_{IDEAL} = \mathbf{x}) = F_{N(0,1)}\left(\frac{\sqrt{|S|}}{\sigma}\right), \quad (1)$$

  where $F_{N(0,1)}$ is the standard-normal distribution function.
- p_{IDEAL} is always below 1, and is increasing in the number of sources $|S|$ and decreasing in the evidence variance σ^2. For instance, it is $p_{IDEAL} \approx 0.868$ if $|S| = 5$ and $\sigma^2 = 4$.

Groups within an infinite population

- Since jury theorems vary the group size n, we straightforwardly extend the simple opinion structure $(\mathbf{x}, (\mathbf{e}_s)_{s\in S}, (\mathbf{S}_i)_{i\in N})$ by letting the set of persons N be the infinite set $\{1, 2, ...\}$, called the 'population'.
- Any group $\{1, ..., n\} \subseteq N$ of a finite size $n \geq 1$ has a majority opinion denoted $\mathbf{o}_{maj,n}$ and a majority competence $p_{maj,n} = Pr(\mathbf{o}_{maj,n} = \mathbf{x})$.

Pre-deliberation jury theorem – stated informally

Theorem: *The group of any size* n

(a) *performs sub-ideally as long as people are not utterly perfect at accessing sources,*

(b) *but reaches the ideal asymptotically as* $n \longrightarrow \infty$ *if people are good enough at accessing sources.*

Pre-deliberation jury theorem – stated formally

Imperfect Access: At least one source $s \in S$ is not surely accessed, i.e., has access probability $p_{i \to s} < 1$ for each person i.

Access Competence: The probability $p_{s \to i}$ that a person $i \in N$ accesses a source $s \in S$ is at least $2^{-1/|S|} + \epsilon$, for some $\epsilon > 0$ independent of i and s.

Pre-Deliberation Jury Theorem: *Given a simple opinion structure* $(\mathbf{x}, (\mathbf{e}_s), (\mathbf{S}_i))$ *for an infinite population, the majority competence* $p_{maj,n}$

(a) *is at most the ideal competence (1), and less than it under Imperfect Access,*

(b) *converges to it as* $n \to \infty$ *under Access Competence.*

Critique

- The assumption of Access Competence (under which group competence is asymptotically ideal) is quite demanding.
- For example, with $|S| = 5$ sources the access probability $p_{s \to i}$ must exceed $2^{-1/5} \approx 0.87$ for all persons i and sources s.
- Fortunately, after deliberation a weaker competence assumption suffices. Why?

Introducing deliberation

- Now suppose the group deliberates prior to voting, following a share-absorb process.
- Formally, consider sharing and absorbing probabilities $\left(p_{s,i\rightarrow}, p_{s,i\leftarrow}\right)_{s\in S, i\in N}$ for the infinite population $N = \{1, 2, ...\}$. For any finite group $\{1, ..., n\} \subseteq N$ (where $n \geq 1$), the parameters $\left(p_{s,i\rightarrow}, p_{s,i\leftarrow}\right)_{s\in S, i\in\{1,...,n\}}$ determine new source sets $\mathbf{S}^{+}_{i,n}$ and competence levels $p^{+}_{i,n}$ of members $i \in \{1, ..., n\}$ and a new group competence $p^{+}_{maj,n}$.

A weaker competence condition

Our Post-deliberation Jury Theorem will need only a weaker competence than Access Competence:

Acquisition Competence: Informally, for all persons i and sources s, the person has a high access probability $p_{s\to i}$ *or* a high absorbing probability $p_{s,i\leftarrow}$ (or both). Formally, for all persons $i \in N$ and sources $s \in S$, the product $(1 - p_{s\to i})(1 - p_{s,i\leftarrow})$ is at most $1 - 2^{-1/|S|} - \epsilon$, for some $\epsilon > 0$ independent of i and s.

Proposition 1 *Acquisition Competence is logically weaker than Access Competence.*

Often, Access Competence fails but Acquisition Competence holds. Intuition: individuals can be bad at accessing, but compensate by being good at absorbing during deliberation.

Post-deliberation jury theorem

Non-Vanishing Participation: For each source $s \in S$, the probability that a person i accesses and shares s, $p_{s\to i} \times p_{s,i\to}$, does not tend to 0 as $i \to \infty$.[6]

Post-Deliberation Jury Theorem: *Given a simple opinion structure* $(\mathbf{x}, (\mathbf{e}_s), (\mathbf{S}_i))$ *and a share-absorb process, both for an infinite population, the post-deliberation majority competence* $p^+_{maj,n}$

(a) *is at most the ideal competence (1), and less than it under Imperfect Access,*

(b) *converges to it as* $n \to \infty$ *under Acquisition Competence and Non-Vanishing Participation.*

[6]This holds for instance if all $p_{s\to i}$ and $p_{s,i\to}$ exceed some fixed level $\epsilon > 0$.

Critique

- By this jury theorem, the interplay of deliberation and group increase makes the group opinion asymptotically ideal under interesting assumptions: individuals must be good enough at acquiring sources ('Acquisition Competence') and not stop participating asymptotically ('Non-Vanishing Participation').
- The comparison of both jury theorems suggests that asymptotically ideal outcomes are easier to achieve with deliberation.

The competence gap

- The difference between ideal and actual group competence, $p_{IDEAL} - p_{maj,n}$, defines the *competence gap*.
- One might try to close it – either through deliberation or through group increase.
- How does each instrument operate? This will soon become clear.

The relatively ideal opinion

- The *relatively ideal opinion* in the group of size n is the opinion based on the available source set $\cup_{i=1}^{n}\mathbf{S}_i$ rather than the full set S, denoted $\mathbf{o}_{ideal,n}$ and defined like $\mathbf{o}_{IDEAL}$ but with 'S' replaced by '$\cup_i\mathbf{S}_i$'.
- Its correctness probability $p_{ideal,n} = Pr(\mathbf{o}_{ideal,n} = \mathbf{x})$ is the *relatively ideal competence*.

Decomposing the competence gap

The competence gap $p_{IDEAL} - p_{maj,n}$ can now be decomposed into the sum of two gaps:

- *Gap 1* is the difference $p_{ideal,n} - p_{maj,n}$ between relatively ideal and actual competence, stemming from imperfect use of available sources.
  –> Deliberation can reduce it!
- *Gap 2* is the difference $p_{IDEAL} - p_{ideal,n}$ between ideal and relatively ideal competence, stemming from the unavailability of some sources in S.
  –> Increasing group size can reduce it!

Illustration

- Example of how both gaps depend on group size and on whether there has been deliberation:

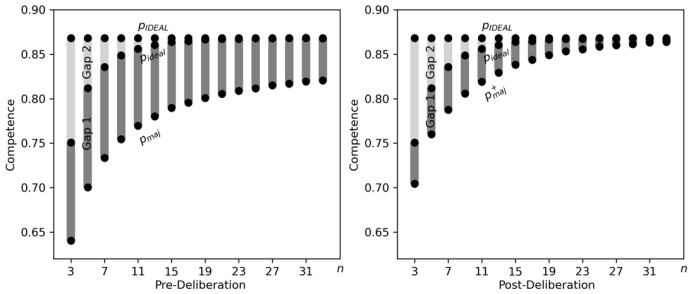

Figure 1: How the competence gaps depend on deliberation and group size

- The parameters were set s.t. Access Competence is violated but Acquisition Competence holds.[7] In result, group competence is asymptotically sub-ideal pre-deliberation, but ideal post-deliberation, in line with our jury theorems.
- Gap 2 is closed asymptotically, i.e., $p_{ideal,n} \longrightarrow p_{IDEAL}$, since larger and larger groups ultimately have all evidence available.

[7]Specifically, $|S| = 5$, $\sigma = 2$, $p_{s\rightarrow i} = 0.2$, $p_{s,i\rightarrow} = 0.5$ and $p_{s,i\leftarrow} = 0.85$. Access Competence is violated because $p_{s\rightarrow i} < 2^{-1/|S|} \approx 0.871$. Acquisition Competence holds because $(1 - p_{s\rightarrow i}) \times (1 - p_{s,i\leftarrow}) = 0.12 < 1 - 2^{-1/|S|} \approx 0.129$.

Part 5

Measuring Failures 1 and 2

Why study failures?

- Jury theorems don't tell *why* deliberation helps or harms.
- Because of raising Failure 1? Or Failure 2?
- To answer these questions, we now present measures of the extent of both failures.

Two types of imbalance

- Failure 1 stems from an imbalance between (the spread of) sources.
- Failure 2 stems from an imbalance between (the strength of evidence of) persons.
- To obtain proxies of Failures 1 and 2, we shall measure both forms of imbalance. How?

Spread imbalance

- The spread of a source s is the number of of its owners $\#\{i : s \in \mathbf{S}_i\}$.
- The spread imbalance between two sources s and s' is the absolute difference in spread, normalised by the average spread.
- The *spread imbalance* simpliciter is the average spread imbalance between distinct sources:

$$\mathbf{SI} = \frac{1}{|S|\left(|S|-1\right)} \sum_{(s,s')\in S^2 : s\neq s'} \text{`spread imbalance between } s \text{ and } s'\text{'}$$

$$= \frac{1}{|S|\left(|S|-1\right)} \sum_{(s,s')\in S^2 : s\neq s'} \frac{|\#\{i : s \in \mathbf{S}_i\} - \#\{i : s' \in \mathbf{S}_i\}|}{\frac{1}{2}\left(\#\{i : s \in \mathbf{S}_i\} + \#\{i : s' \in \mathbf{S}_i\}\right)}.$$

Interpersonal imbalance

- The evidence strength of a person i is $\left|\sum_{s \in \mathbf{S}_i} \mathbf{e}_s\right|$.
- The interpersonal imbalance between two individuals i and j is the change in evidence strength, normalised by the average evidence strength.
- The *interpersonal imbalance* simpliciter is the average imbalance between pairs of individuals:

$$\begin{aligned}\mathbf{II} &= \frac{1}{n(n-1)} \sum_{(i,j) \in N^2 : i \neq j} \text{‘imbalance between } i \text{ and } j\text{’} \\ &= \frac{1}{n(n-1)} \sum_{(i,j) \in N^2 : i \neq j} \frac{\left| \left|\sum_{s \in \mathbf{S}_i} \mathbf{e}_s\right| - \left|\sum_{s \in \mathbf{S}_j} \mathbf{e}_s\right| \right|}{\frac{1}{2}\left(\left|\sum_{s \in \mathbf{S}_i} \mathbf{e}_s\right| + \left|\sum_{s \in \mathbf{S}_j} \mathbf{e}_s\right|\right)}.\end{aligned}$$

Resulting vs. systemic imbalance

- The indices $\mathbf{SI}$ and $\mathbf{II}$ measure *resulting imbalance*: they depend on the (random) nature of people's sources and evidences
- *Systemic imbalance* consists in a tendency towards (resulting) imbalance, measured by expected resulting imbalance.
- *Notation:*
  - Systemic source imbalance: $\mathcal{SI} = \mathbb{E}(\mathbf{SI})$
  - Systemic interpersonal imbalance: $\mathcal{II} = \mathbb{E}(\mathbf{II})$
- *Terminology*: We use the term 'imbalance' in both senses, as the context makes clear which type is meant.
- The systemic indices $\mathcal{SI}$ and $\mathcal{II}$ will be our failure proxies.

Part 6

A typology of beneficial and harmful deliberation

Goal

- Use Monte-Carlo simulations to compare
  - old and new group competence p_{maj} and p^{+}_{maj},
  - old and new (proxy of) Failure 1 $\mathcal{SI}$ and $\mathcal{SI}^{+}$,
  - old and new (proxy of) Failure 2 $\mathcal{II}$ and $\mathcal{II}^{+}$.
- Do this for many 'scanners' (= parameter constellations),
- Identify a typology of 'good' and 'bad' forms of deliberation.

Working assumptions in simulations

- Simple opinion structure, share-absorb deliberation (genuine restrictions!)
- 100,000 simulation rounds (enough for reliable estimates)
- $n = 9$, $|S| = 5$, $\sigma = 2$ (not much hinges on these values, by our robustness checks)

Even vs. equal deliberation

We call a share-absorb process:

- *even* if its parameters are source-independent. Intuitively, sources are treated symmetrically.
- *equal* if its parameters are person-independent. Intuitively, everyone takes part equally in deliberation.

The labels ‘even’ and ‘equal’ can be also applied to sharing alone, or to absorbing alone, or to access

Participatory deliberation

We call a share-absorb process:

- *participatory* if every person i shares substantially, in the sense that the average sharing probability $\frac{1}{|S|}\sum_{s\in S} p_{s,i\rightarrow}$ exceeds some threshold δ (of for instance 0.5). For even deliberation this condition simply means that everyone's (source-independent) sharing probability exceeds δ. There are stronger and weaker notions of 'participatory', depending on δ.

Even, equal and participatory deliberation is beneficial

	Parameters			Pre-Deliberation			Post-Deliberation			change from...		
#	$p_{s\to i}$	$p_{s,i\to}$	$p_{s,i\leftarrow}$	p_{maj}	$\mathcal{SI}$	$\mathcal{II}$	p^+_{maj}	$\mathcal{SI}^+$	$\mathcal{II}^+$	p_{maj} to p^+_{maj}	$\mathcal{SI}$ to $\mathcal{SI}^+$ in %	$\mathcal{II}$ to $\mathcal{II}^+$ in %
1.1	0.5	0.5	0.5	.839	.394	.842	.859	.300	.602	.020	-24.0	-28.5
1.2	0.8	0.8	0.8	.865	.189	.528	.868	.064	.166	.003	-66.3	-68.5
1.3	0.5	0.8	0.8	.839	.395	.842	.867	.146	.344	.028	-63.0	-59.1
1.4	0.2	0.5	0.5	.755	.847	1.239	.800	.953	.825	.045	12.5	-33.4

Figure 2: Results for even, equal and participatory deliberation.

- Deliberation always raises maj. competence and lowers Failure 2.
- It sometimes raises Failure 1, but this effect never dominates.

Five harmful types of deliberation

- We found 5 elementary types of harmful deliberation; these types and some 'hybrid' types that combine them seem to exhaust the space of harmful share-absorb processes (except for highly degenerate access, e.g., when everyone surely accesses all sources).
- Figure **??** gives an example of each elementary type.
- N.B.: Deliberation always harms by raising Failure 1, not 2.

	Parameters			**Pre-Deliberation**			**Post-Deliberation**			**change from...**		
#	$p_{s\to i}$	$p_{s,i\to}$	$p_{s,i\leftarrow}$	p_{maj}	$\mathcal{SI}$	$\mathcal{II}$	p^+_{maj}	$\mathcal{SI}^+$	$\mathcal{II}^+$	p_{maj} to p^+_{maj}	$\mathcal{SI}$ to $\mathcal{SI}^+$ in %	$\mathcal{II}$ to $\mathcal{II}^+$ in %
	Fully private evidence											
2.1	1|0	0.5	0.5	.890	0	.879	.874	.726	.780	-.016	∞	-11.3
	Non-participatory deliberation											
2.2	0.2	0.1	1	.754	.847	1.239	.749	1.013	.822	-.005	19.5	-33.6
	Uneven sharing 1:4											
2.3	0.2	1|0	1	.755	.847	1.239	.740	1.077	.703	-.015	27.1	-43.3
	Uneven absorbing 1:4											
2.4	0.2	1	1|0	.755	.847	1.239	.741	1.077	0.702	-.015	27.2	-43.3
	Unequal sharing 1:8											
2.5	0.2	1|0	1	.756	.847	1.239	.748	1.033	.758	-.007	22.1	-38.8

Harmful deliberation of type 1: some private-evidence scenarios

- Here most or all members have few or no evidences in common: their source sets have little or no overlap.
- Example: Scenarios 2.1.
  - N.B.: Its access structure is implicit in Condorcet's jury theorem, and essentially equivalent to Austen-Smith and Banks' (1996) model.
- *The harmful mechanism:*
  - The group is already highly competent ex ante (low Failure 1, high voter independence)
  - Deliberation unsettles this fine balance, creating Failure 1 and voter dependence.
- N.B.: Even if deliberation is perfectly even and equal, as in Scenario 2.1

Harmful deliberation of type 2: some private-evidence scenarios

- Here many voters have low average sharing probability.
- Example: Scenario 2.2. Sharing probability only 0.1.
- *The harmful mechanism:* By the combination of low sharing and high absorbing, deliberation puts very few evidences on the table (e.g., only source s_3 in Figure 3), which are widely absorbed.

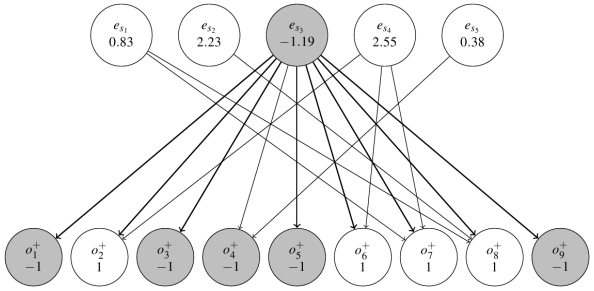

Figure 3: An epistemically harmful outcome of even deliberation. Thin arrows: initial access. Thick arrows additional post-deliberation access. Evidence values and post-deliberation opinions also displayed. Correct option: 1 (so source s_3 supports incorrect option).

Harmful deliberation of types 3 & 4: some uneven-sharing or uneven-absorbing scenarios

- Here sources are shared unevenly (example: Scenario 2.3) or absorbed unevenly (example: Scenario 2.4).
- *The harmful mechanism:* Certain evidences are singled out for wide spread, raising Failure 1
- The effect is at its worst if very few or just one evidence is put on the table (Scenario 2.3) or picked up (Scenario 2.4).

Harmful deliberation of types 5: some unequal-sharing scenarios

- Here some members share far more actively than others.
- Example: Scenario 2.5, where one member shares everything, the others nothing.
- *The harmful mechanism:* Again, certain evidences spread very widely, raising Failure 1.
- Interestingly, even though this mechanism operates through person- rather than source-dependence, it is yet again Failure 1 (not 2) that rises.

Recommendation

- Our analysis yields a recommendation: the group should engage in *participatory and even* deliberation.
  - So: source-independent sharing and absorbing probabilities $p_{s,i\rightarrow} \equiv p_{i\rightarrow}$ and $p_{s,i\leftarrow} \equiv p_{i\leftarrow}$ ('even') and sufficiently high $p_{i\rightarrow}$ ('participatory').
- **Reason**: This the 5 harmful types of scenario.
  - Type 2 is blocked as it is non-participatory.
  - Types 3 & 4 are blocked as they are uneven.
  - Types 1 and 5 are blocked because, though they can be even, they never are very participatory (e.g., the sharing probability is only $\frac{1}{2}$ in 2.1 and only $\frac{1}{9}$ on average in 2.5).

Part 7

Generalised opinion structures and deliberation processes

Generalising opinion structures

Two formal generalisations:

- Drop all distributional restrictions, i.e., Equiprobable States, Simple Gaussian Evidences, and Independent Sources.
- Drop the presupposition that each person aggregates their evidence *additively* (which becomes ad hoc for general evidence distributions).

Beyond evidence

- A variable $\mathbf{e}_s$ cannot generally be called an 'evidence', as it can be independent of the state.
- We call it an 'influences', which we further call
  - an *evidence* if it is dependent of the state,
  - a *noise* if it is independent of the state.
- Real people are influenced by noises!

In sum

A *general opinion structure* is a quadruple $(\mathbf{x}, (\mathbf{e}_s)_{s\in S}, (\mathbf{S}_i)_{i\in N}, f) \equiv (\mathbf{x}, (\mathbf{e}_s), (\mathbf{S}_i), f)$ of:

(1) a random variable $\mathbf{x}$, the *state* or *correct option*, taking the value $\mathbf{1}$ or $-\mathbf{1}$ with arbitrary (non-zero) probabilities;

(2) a family $(\mathbf{e}_s)$, indexed by some set S of *sources* (non-empty and finite), consisting of real-valued random variables, the *influences* from these sources, with arbitrary (discrete or continuous) distributions and interdependencies;

(3) a family $(\mathbf{S}_i)$, indexed by some set N of *persons*, consisting of random subsets of S, the *source sets* of these persons, again with arbitrary distributions and interdependencies;

(4) An *influence aggregator* f, mapping any influence bundle $(e_s)_{s\in S'}$ ($S' \subseteq S$) to an 'aggregate influence' $f((e_s)_{s\in S'}) \in \mathbb{R}$.

Simple opinion structures as a special case

A simple opinion structure $(\mathbf{x}, (\mathbf{e}_s)_{s\in S}, (\mathbf{S}_i))$ is a special case, characterised by several distributional restrictions and by an (implicit) *additive* influence aggregator f given by

$$f((e_s)_{s\in S'}) = \sum_{s\in S'} e_s.$$

By-products

- Our entire machinery carries over to a general opinion structure $(\mathbf{x}, (\mathbf{e}_s), (\mathbf{S}_i), f)$.
- The *opinion* of person i is determined by her (now possibly non-additive) aggregate influence:

$$\mathbf{o}_i = \begin{cases} 1 & \text{if } f((\mathbf{e}_s)_{s \in \mathbf{S}_i}) > 0 \\ -1 & \text{if } f((\mathbf{e}_s)_{s \in \mathbf{S}_i}) < 0 \\ 0 & \text{if } f((\mathbf{e}_s)_{s \in \mathbf{S}_i}) = 0. \end{cases}$$

- All other derivative concepts – notably personal competence p_i, the majority opinion $\mathbf{o}_{maj}$, majority competence p_{maj}, and spread imbalance $\mathbf{SI}$ – keep their original definitions, except from interpersonal imbalance $\mathbf{II}$, whose definition should of course be generalised by aggregating someone's influences using f rather than simple summation.

Discussion of general opinion structures

- **Evidential versus non-evidential influences.**
- **Discrete versus continuous influences.**

General deliberation processes

- Share-absorb processes are just one type of process generating a new source profile ($\mathbf{S}_i^+$).
- Deterministic deliberation processes: cf. paper
- General definition: cf. paper
- Examples: cf. paper

Part 8

A theoretic analysis of Failure 3