Aggregating Judgements: Logical and Probabilistic Approaches

Lecture 5

Eric Pacuit
Department of Philosophy
University of Maryland
pacuit.org

August 10, 2018

Plan

- ✓ **Monday** Representing judgements; Introduction to judgement aggregation; Aggregation paradoxes I
- ✓ **Tuesday** Aggregation paradoxes II, Axiomatic characterizations of aggregation methods I
- ✓ **Wednesday** Axiomatic characterizations of probabilistic opinions
- ✓ **Thursday** Pooling imprecise probabilities; Distance-based characterizations; Aumann's agreeing to disagree theorem and related results
- **Friday** Merging of probabilistic opinions (Blackwell-Dubins Theorem); Belief polarization; Concluding remarks

Learning in a group

1. Start with the same beliefs, receive the same evidence. (Convergence)
2. Start with the same beliefs, receive different evidence.
3. Start with different beliefs, receive the same evidence.
4. Start with different beliefs, receive different evidence. (Polarization)

Events

Suppose that W is a non-empty set and $\mathcal{F}$ is a σ-field on W.

The elements of W can be thought of as possible worlds and the members of $\mathcal{F}$ as propositions.

Events

Suppose that W is a non-empty set and $\mathcal{F}$ is a σ-field on W.

The elements of W can be thought of as possible worlds and the members of $\mathcal{F}$ as propositions.

Example: W is set of all infinite sequences of coin tosses ($W = \{\sigma \mid \sigma \in \{H, T\}^\omega\}$ and $\mathcal{F}$ is all propositions about coin tossing events of interest:

- $[H_n = k]$: There are k heads in the first n rounds.
- $[\lim_{n\to\infty} H_n = 0.5]$: The number of heads *converges* to 0.5 as the number of flips increases (e.g., the coin is unbiased).

Evidence

Let $\mathcal{E}_1, \mathcal{E}_2, \ldots, \mathcal{E}_n, \ldots$ be an infinite sequence of partitions on W such that $\mathcal{E}_{n+1}$ *refines* $\mathcal{E}_n$.

$\mathcal{E}_n$ is the information that the agent receives at time n.

$\mathcal{E}_n[w]$ is the element of $\mathcal{E}_n$ containing w.

For each n, let $\mathcal{F}_n$ be the σ-algebra generated by $\mathcal{E}_n$. We assume that $\mathcal{F} = \cup_n \mathcal{F}_n$.

Example

$$W = \{\sigma \mid \sigma \in \{H, T\}^\omega\}$$

Notation: X^ω is the set of infinite strings over X;
$\sigma \sqsubseteq \sigma'$ means σ is an initial segment of σ'.
$[\sigma] = \{\sigma' \mid \sigma \sqsubseteq \sigma'\}$, e.g., $[H] = \{\sigma \mid H \sqsubseteq \sigma\}$

Example

$$W = \{\sigma \mid \sigma \in \{H, T\}^{\omega}\}$$

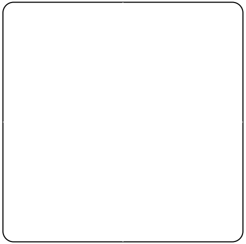

1. $\mathcal{E}_0 \quad = \quad \{W\}$

Notation: X^{ω} is the set of infinite strings over X;
$\sigma \sqsubseteq \sigma'$ means σ is an initial segment of σ'.
$[\sigma] = \{\sigma' \mid \sigma \sqsubseteq \sigma'\}$, e.g., $[H] = \{\sigma \mid H \sqsubseteq \sigma\}$

Example

$$W = \{\sigma \mid \sigma \in \{H, T\}^\omega\}$$

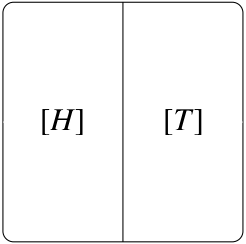

1. $\mathcal{E}_0 \quad = \quad \{W\}$
2. $\mathcal{E}_1 \quad = \quad \{[H], [T]\}$

Notation: X^ω is the set of infinite strings over X;
$\sigma \sqsubseteq \sigma'$ means σ is an initial segment of σ'.
$[\sigma] = \{\sigma' \mid \sigma \sqsubseteq \sigma'\}$, e.g., $[H] = \{\sigma \mid H \sqsubseteq \sigma\}$

Example

$$W = \{\sigma \mid \sigma \in \{H, T\}^\omega\}$$

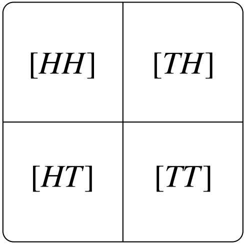

1. $\mathcal{E}_0 = \{W\}$
2. $\mathcal{E}_1 = \{[H], [T]\}$
3. $\mathcal{E}_2 = \{[HH], [HT], [TH], [TT]\}$

$$[H] = [HH] \cup [HT] \qquad [T] = [TH] \cup [TT]$$

Notation: X^ω is the set of infinite strings over X;
$\sigma \sqsubseteq \sigma'$ means σ is an initial segment of σ'.
$[\sigma] = \{\sigma' \mid \sigma \sqsubseteq \sigma'\}$, e.g., $[H] = \{\sigma \mid H \sqsubseteq \sigma\}$

Example

$$W = \{\sigma \mid \sigma \in \{H, T\}^\omega\}$$

$[HH]$	$[TH]$
$[HT]$	$[TT]$

1. $\mathcal{E}_0 = \{W\}$
2. $\mathcal{E}_1 = \{[H], [T]\}$
3. $\mathcal{E}_2 = \{[HH], [HT], [TH], [TT]\}$

$\vdots \quad \vdots \quad \vdots$

$$[H] = [HH] \cup [HT] \qquad [T] = [TH] \cup [TT]$$

Notation: X^ω is the set of infinite strings over X;
$\sigma \sqsubseteq \sigma'$ means σ is an initial segment of σ'.
$[\sigma] = \{\sigma' \mid \sigma \sqsubseteq \sigma'\}$, e.g., $[H] = \{\sigma \mid H \sqsubseteq \sigma\}$

Learning

For $A \in \mathcal{F}$, $P(A \mid \mathcal{E}_n[w]) = \frac{P(A \cap \mathcal{E}_n[w])}{P(\mathcal{E}_n[w])}$

Convergence to Certainty

$$P(A \mid \mathcal{E}_n[w]) \to \begin{cases} 1 & w \in A \\ 0 & w \notin A \end{cases}$$

for **almost every** w with regard to the prior probability P.

The set of w for which the above does not hold has measure 0.

Convergence to certainty yields a first pass on merging of opinions.

Suppose that Ann's degrees of beliefs are represented by P and Bob's by Q, and let $P_n[A](w) = P(A \mid \mathcal{E}_n[w])$ and $Q_n[A](w) = Q(A \mid \mathcal{E}_n[w])$.

Then $P_n[A]$ and $Q_n[A]$ both converge to zero or to one almost surely with respect to the priors P and Q, respectively.

Now, Ann believes with certainty that $P_n[A]$ and $Q_n[A]$ will agree in the limit ***whenever she assigns probability one to any set to which Bob assigns probability one***.

Now, Ann believes with certainty that $P_n[A]$ and $Q_n[A]$ will agree in the limit ***whenever she assigns probability one to any set to which Bob assigns probability one***.

This result applies to particular events A. But it does not say anything about Ann's and Bob's overall conditional probabilities.

The Blackwell-Dubins theorem fills this gap.

D. Blackwell and L. Dubins. *Merging of opinions with increasing information*. The Annals of Mathematical Statistics, 33, pp. 882 - 886.

The Blackwell-Dubins theorem fills this gap.

D. Blackwell and L. Dubins. *Merging of opinions with increasing information*. The Annals of Mathematical Statistics, 33, pp. 882 - 886.

- As Ann and Bob observe events and update by Bayesian conditioning, their probabilities become close *uniformly* in all events.

The Blackwell-Dubins theorem fills this gap.

D. Blackwell and L. Dubins. *Merging of opinions with increasing information*. The Annals of Mathematical Statistics, 33, pp. 882 - 886.

- As Ann and Bob observe events and update by Bayesian conditioning, their probabilities become close *uniformly* in all events.

- The Blackwell-Dubins theorem does not require that probabilities converge (as in convergence to certainty): Ann's and Bob's conditional probabilities may get closer even if they don't converge.

Variational Distance

Suppose that μ and ν are two measures over all events in $\mathcal{F}$.

$$d(\mu, \nu) = \sup_{A \in \mathcal{F}} |\mu(A) - \nu(A)|$$

Variational Distance

Suppose that μ and ν are two measures over all events in $\mathcal{F}$.

$$d(\mu, \nu) = \sup_{A \in \mathcal{F}} |\mu(A) - \nu(A)|$$

P is said to **merge to** Q if for Q almost every w, $d(P_n(w), Q_n(w)) \to 0$ as $n \to \infty$.

Absolute Continuity

Q is absolutely continuous relative to P ($Q \ll P$) if for all $A \in \mathcal{F}$,

$$Q(A) > 0 \implies P(A) > 0$$

Blackwell-Dubins Theorem. If $Q \ll P$, then P merges to Q.

D. Blackwell and L. Dubins. *Merging of opinions with increasing information*. The Annals of Mathematical Statistics, 33, pp. 882 - 886.

- The Blackwell-Dubins Theorem assumes that the agents agree about which events are certain.

- The Blackwell-Dubins Theorem assumes that the agents agree about which events are certain.

- In the Blackwell-Dubins Theorem, the agents receive *hard evidence* and update by conditioning.

- The Blackwell-Dubins Theorem assumes that the agents agree about which events are certain.

- In the Blackwell-Dubins Theorem, the agents receive *hard evidence* and update by conditioning. What about...
  - learning from others?
  - alternative update rules?
  - ambiguous evidence?

Deliberation

In an ideal situation, the discussion will elicit from each member of the group not only their judgements, but also their reasons, arguments and evidence that back up these judgements. Through discussion and debate, the group can sort through all of the evidence and arguments leading to a more informed solution.

Deliberation

In an ideal situation, the discussion will elicit from each member of the group not only their judgements, but also their reasons, arguments and evidence that back up these judgements. Through discussion and debate, the group can sort through all of the evidence and arguments leading to a more informed solution.

A common criticism of unstructured group discussion is that it enhances cognitive errors rather than mitigates them.

Deliberation - Problems

Bias against the minority: There is a tendency for groups to ignore isolated, minority or lower-status members.

Deliberation - Problems

Bias against the minority: There is a tendency for groups to ignore isolated, minority or lower-status members.

Anchoring effect: There is a tendency to rely too heavily, or “anchor”, a judgement on one piece of information (for example, the first announced judgement, the judgement of the most senior person in the group, or the judgement of the loudest person in the group).

Deliberation - Problems

Bias against the minority: There is a tendency for groups to ignore isolated, minority or lower-status members.

Anchoring effect: There is a tendency to rely too heavily, or "anchor", a judgement on one piece of information (for example, the first announced judgement, the judgement of the most senior person in the group, or the judgement of the loudest person in the group).

Common knowledge effect: Information held by all members of the group has more influence on the final decision than information held by only a few members of the group. So, if everybody in the group has some information, then it is more valuable than the information of just a few members.

Deliberation - Problems

Suppose that there are 10 people estimating a parameter whose true value is 40 with the following initial estimates:

1	2	3	4	5	6	7	8	9	10	Avg
15	18	20	22	30	45	50	55	60	61	37.6

Deliberation - Problems

Suppose that there are 10 people estimating a parameter whose true value is 40 with the following initial estimates:

1	2	3	4	5	6	7	8	9	10	Avg
15	18	20	22	30	45	50	55	60	61	37.6
16	25	21	23	31	41	41	40	41	45	32.4

M. Burgman, M. McBride, R. Ashton, A. Speirs-Bridge, L. Flander, B. Wintle, F. Fidler, L. Rumpff, and C. Twardy. *Expert Status and Performance*. PLoS One 6(7), 2011.

Belief Polarization

A. Bramson, P. Grim, D. J. Singer, W. J. Berger, G. Sack, S. Fisher, C. Flocken, and B. Holman (2017). *Understanding Polarization: Meanings, Measures, and Model Evaluation*. Philosophy of Science, 84 (1), 115?159.

The term ‘polarization’ is sometimes used to refer to a static property (the population is polarized) and sometimes a process (the population is polarizing)

Spread: breadth of opinions: how far apart are the extremes?

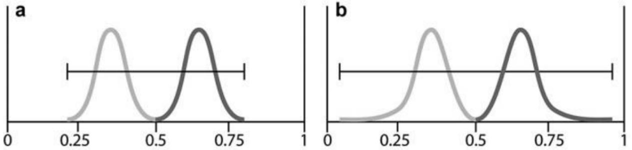

Figure 1. Belief distribution *b* shows greater polarization in the sense of spread than does belief distribution *a*. Two separate groups are shown, but that is irrelevant to polarization in the sense of spread. Color version available as an online enhancement.

Dispersion: Statistical dispersion (or statistical variation). Any of various measures of statistical dispersion might be used: mean difference, average absolute deviation, standard deviation, coefficient of variation, or entropy.

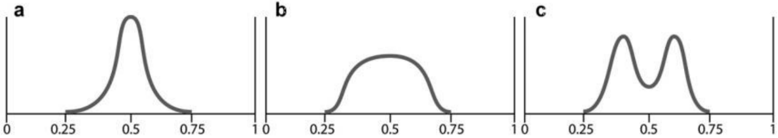

Figure 2. Distribution *c* shows greater polarization in the sense of dispersion than does belief distribution *b*, which is greater than distribution *a*. Color version available as an online enhancement.

Coverage: A polarized group is thought of as one with little diversity of opinion, one in which only narrow bands of the opinion space are occupied.

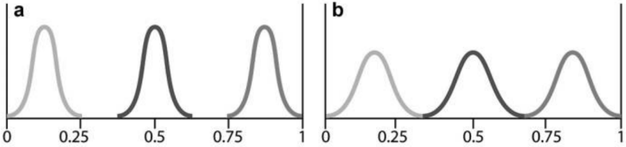

Figure 3. Distribution *a* is more polarized than *b* in the sense of representing less coverage on the spectrum of potential belief. Although the plots show groups and differences in heights, neither of those features are aspects of polarization in the sense of coverage. Color version available as an online enhancement.

Distinctness: The degree to which the group distributions can be separated.

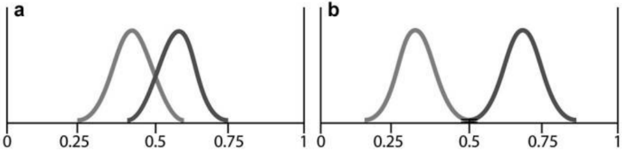

Figure 8. Distribution *b* shows greater polarization than *a* in terms of distinctness. Color version available as an online enhancement.

Learning from others

If Q reports that the probability of E is q, then $P_{new}(E) = P(E \mid Q(E) = q)$

Learning from others

If Q reports that the probability of E is q, then $P_{new}(E) = P(E \mid Q(E) = q)$

Shortcut: $P_{new}(E) = \frac{P(E)+Q(E)}{2}$

C. Genest and M. J. Schervish (1985). *Modeling expert judgments for bayesian updating*. The Annals of Statistics, 13(3), pgs. 1198?1212.

R. Bradley (2017). *Learning from others: conditioning versus averaging*. Theory and Decision.

J. W. Romeijn and O. Roy (2017). *All agreed: Aumann meets DeGroot*. Theory and Decision.

The DeGroot/Lehrer Model

- There is a proposition about which several individuals disagree.
- Each individual i initially assigns some probability p_i to the proposition.

The DeGroot/Lehrer Model

- There is a proposition about which several individuals disagree.
- Each individual i initially assigns some probability p_i to the proposition.

- Each individual i assigns every individual j (including himself!) a non-zero weight $w_{i,j}$. The weights represent how reliable i believes j is relative to others in the group.
  - $0 \leq w_{i,j} \leq 1$
  - For any individual, the weights sum to 1, i.e., $\sum_{i,j} w_{i,j} = 1$

The DeGroot/Lehrer Model

- Time is divided into discrete stages.
- Let i's degree of belief on stage t be represented by $p_{i,t}$ At stage $t + 1$, individual i updates his belief to be a weighted-average of everyone's beliefs from stage t.

$$p_{i,t+1} = \sum_j w_{i,j} p_{j,t}$$

The DeGroot/Lehrer Model

Theorem. [DeGroot, 1974, Lehrer and Wagner, 1981] In the above model, all individuals beliefs approach a common probability as the number of stages grows larger.

The DeGroot/Lehrer Model

- Why should individuals assign non-zero weight to others?
- Why should individuals repeat the averaging process?
- Why should the weights remain constant?

Repeated Averaging

> *[R]efusing to shift from state 1 to state 2 is equivalent to assigning a weight of 0 to other members of the group at this stage. This amounts to the assumption that there is no chance that one is mistaken and no chance that others in the group with whom one disagrees are correct. In short, the only alternative to the iterated aggregation converging toward a consensual probability assignment is individual dogmatism at some stage.*

(Lehrer, 1976, pg. 331)

Hegselmann and Krause

An individual's peers change over time.

In the model, an individual considers only those whose opinions are sufficiently similar to his or her own.

Hegselmann and Krause wanted to explain why groups might become polarized and not reach a consensus.

R. Hegselmann and U. Krause (2002). *Opinion dynamics and bounded confidence models, analysis, and simulation*. Journal of artificial societies and social simulation, 5 (3).

The Hegselmann-Krause Model

Agent i's belief is represented by a real number r_i

- A subjective probability of a proposition
- A numerical estimate of some quantity

The truth is represented by a real-number T.

- 0 or 1 might represent the truth-value of some proposition.
- The actual value of some quantity

The Hegselmann-Krause Parameters

- **Trust**: A number ϵ (for all agents) between 0 and 1 that represents how "close" other's opinions must be to one's own in order for one to take them seriously.
- **Truth Seeking**: A number τ (for all agents) between 0 and 1 that represents how strongly she is "attracted" to the truth.
  - $\tau = 0$ means that the agents only listens to their peers.
  - $\tau = 1$ means the agents have direct access to the truth.
- **Time**: Discrete stages $1, 2, 3, \ldots$

The Hegselmann-Krause Model

Let $b_{i,t}$ be i's belief at time t

Let $N_\rho(i,t)$ be those peers whose opinions differ from i's by no more than ρ at stage t:

$$N_\rho(i,t) = \{j \mid |b_{j,t} - b_{i,t}| < \rho\}$$

Let $N = |N_\rho(i,t)|$. Then,

$$b_{i,t+1} = \tau \cdot T + (1-\tau) \sum_{j \in N_\rho(i,t)} \frac{1}{N} \cdot b_{j,t}$$

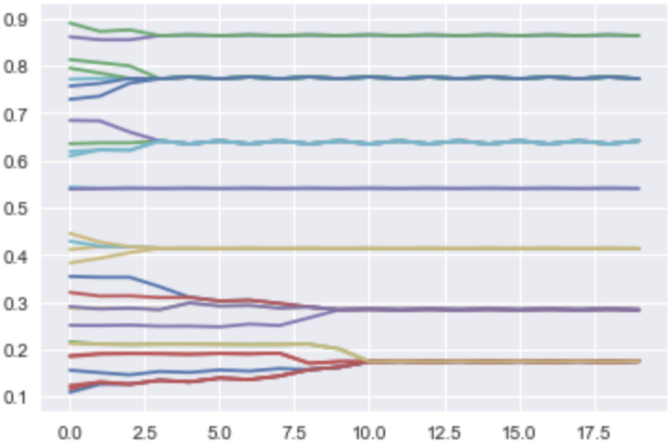

$N = 30 \qquad Rounds = 20 \qquad \tau = 0 \qquad \rho = 0.05$

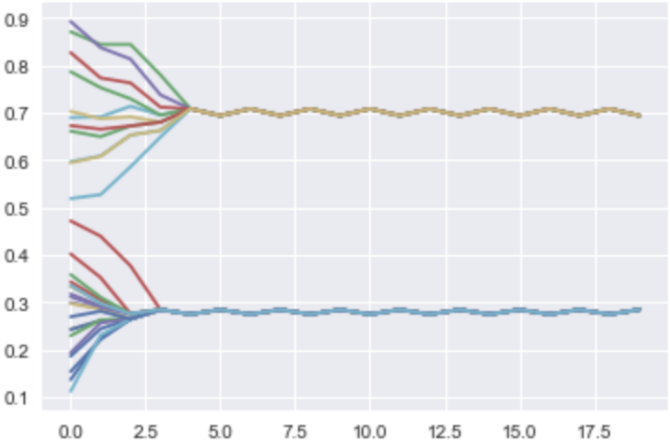

$N = 30 \qquad Rounds = 20 \qquad \tau = 0 \qquad \rho = 0.15$

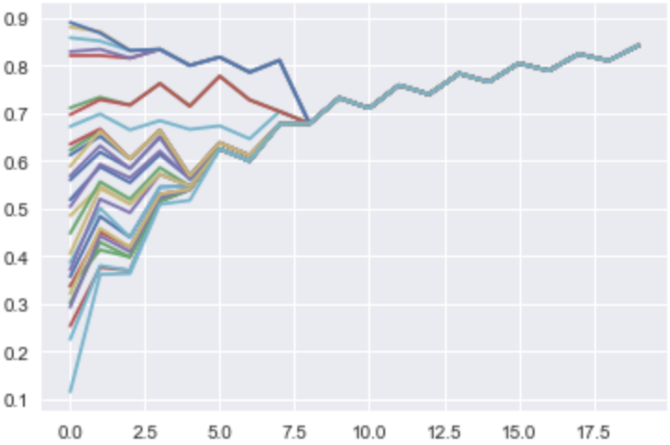

$N = 30 \qquad Rounds = 20 \qquad \tau = 0.1 \qquad \rho = 0.15$

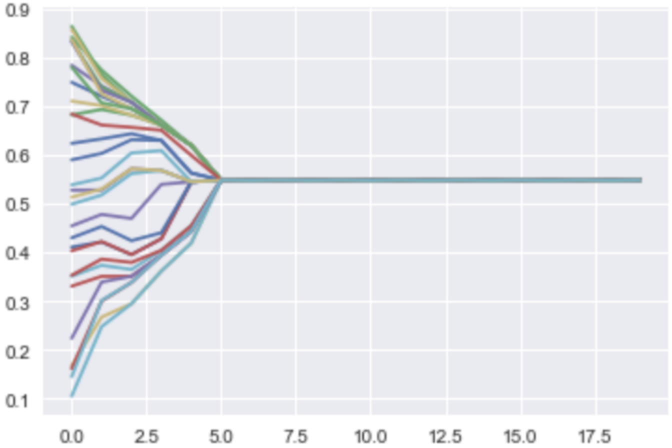

$N = 30 \qquad Rounds = 20 \qquad \tau = 0 \qquad \rho = 0.25$

Epistemic Networks

K. Zollman, Kevin (2013). *Network epistemology: Communication in epistemic communities*. Philosophy Compass, 8 (1), 15 - 27.

V. Bala and S. Goyal (1998). *Learning from neighbors*. Review of Economic Studies, 65 (3), 595 - 621.

B. Golub and M. O. Jackson (2010). *Naive learning in social networks and the wisdom of crowds*. American Economic Journal: Microeconomics, pages 112 - 149.

Epistemic Networks

K. Zollman, Kevin (2013). *Network epistemology: Communication in epistemic communities*. Philosophy Compass, 8 (1), 15 - 27.

V. Bala and S. Goyal (1998). *Learning from neighbors*. Review of Economic Studies, 65 (3), 595 - 621.

B. Golub and M. O. Jackson (2010). *Naive learning in social networks and the wisdom of crowds*. American Economic Journal: Microeconomics, pages 112 - 149.

(only a partial list...)

	w_1	w_2	w_3	w_4
P_1	1/6	1/4	1/3	1/4
P_2	1/2	1/12	1/4	1/6

$\overset{Learn(E)}{\Longrightarrow}$

	w_1	w_2	w_3	w_4
P_1^E	1/3	0	2/3	0
P_2^E	2/3	0	1/3	0

$E = \{w_1, w_3\}$
$A = \{w_1, w_2\}$

	w_1	w_2	w_3	w_4
P_1	$1/6$	$1/4$	$1/3$	$1/4$
P_2	$1/2$	$1/12$	$1/4$	$1/6$

$\overset{Learn(E)}{\Longrightarrow}$

	w_1	w_2	w_3	w_4
P_1^E	$1/3$	0	$2/3$	0
P_2^E	$2/3$	0	$1/3$	0

$E = \{w_1, w_3\}$
$A = \{w_1, w_2\}$

Learning E makes the agents polarized about A:

$P_1(A \mid E) = 1/3 < 5/12 = P_1(A) \leq P_2(A) = 7/12 < 2/3 = P_2(A \mid E)$

T. Herron, T. Seidenfeld and L. Wasserman. *Divisive conditioning: Further results on dilation*. Philosophy of Science 64 (3), pp. 411 - 444, 1997.

A. P. Pedersen and G. Wheeler. *Dilation, disintegrations, and delayed decisions*. In ISIPTA15: Proceedings of the 9th International Symposium on Imprecise Probability: Theories and Applications, 2015.

	G	H	$G \cap H$	$A = (G \cap H) \cup (G^c \cap H^c)$
P_1	0.1	0.5	$P_1(G) \cdot P_1(H)$	0.5
P_2	0.9	0.5	$P_1(G) \cdot P_1(H)$	0.5

	G	H	$G \cap H$	$A = (G \cap H) \cup (G^c \cap H^c)$
P_1	0.1	0.5	$P_1(G) \cdot P_1(H)$	0.5
P_2	0.9	0.5	$P_1(G) \cdot P_1(H)$	0.5

Learn H

$$P_i(A \mid H) = \frac{P_i(((G \cap H) \cup (G^c \cap H^c)) \cap H)}{P_i(H)} = \frac{P_i(G \cap H)}{P_i(H)} = P_i(G)$$

	G	H	$G \cap H$	$A = (G \cap H) \cup (G^c \cap H^c)$
P_1	0.1	0.5	$P_1(G) \cdot P_1(H)$	0.5
P_2	0.9	0.5	$P_1(G) \cdot P_1(H)$	0.5

Learn H

$$P_i(A \mid H) = \frac{P_i(((G \cap H) \cup (G^c \cap H^c)) \cap H)}{P_i(H)} = \frac{P_i(G \cap H)}{P_i(H)} = P_i(G)$$

	G	H	$G \cap H$	$A = (G \cap H) \cup (G^c \cap H^c)$
P_1^H	0.1	1	$P_1(G) \cdot P_1(H)$	0.1
P_2^H	0.9	1	$P_1(G) \cdot P_1(H)$	0.9

S. Huttegger. *Merging of Opinions and Probability Kinematics*. Review of Symbolic Logic, 2015.

Jeffrey Conditionalization

When an observation bears directly on the probabilities over a partition $\{E_i\}$, changing them from $P(E_i)$ to $Q(E_i)$, the new probability Q for any proposition H should be

$$Q(H) = \sum_i P(H \mid E_i) q(E_i)$$

Jeffrey Conditionalization

When an observation bears directly on the probabilities over a partition $\{E_i\}$, changing them from $P(E_i)$ to $Q(E_i)$, the new probability Q for any proposition H should be

$$Q(H) = \sum_i P(H \mid E_i)q(E_i)$$

Fact: If Q is obtained from p by Jeffrey Conditioning on the partition $\{E, \overline{E}\}$ with $Q(E) = 1$, then $Q(\cdot) = P(\cdot \mid E)$.

Agents assumes to undergo learning experiences that yield a sequence $P_1, P_2, \ldots$ of probability measures on $(W, \mathcal{F}_1), (W, \mathcal{F}_2), \ldots$.

Each agent believes with probability one that she will revise her probabilities by performing *probability kinematics* with $p_1, p_2, \ldots$. Each probability measure P_n is fully determined by attaching probability values to members of $\mathcal{E}_n$.

$$P_n(A) = \sum_{E \in \mathcal{E}_n} P(A \mid E) p_n(E)$$

Suppose now that there are two prior probability measures P and Q which are updated successively by probability kinematics on $\mathcal{E}_1, \mathcal{E}_2, \ldots$ using the distributions $p_1, p_2, \ldots$ and $q_1, \ldots$, respectively.

Using the Jeffrey update rule this leads to the new probability measures P_n and Q_n on $\mathcal{F}$ for $n \geq 1$.

It is quite obvious that arbitrary choices of sequences $p_1, p_2, \ldots$ and $q_1, q_2, \ldots$ need not lead to merging. But this is also true for conditioning.

Recall that one requirement of the Blackwell-Dubins theorem is that agents condition on the *same factual evidence*.

Thus, the important question is whether beliefs merge for probability kinematics whenever p_n and q_n represent the same uncertain information.

But what does it mean to get the same uncertain evidence?

Hard Jeffrey Shift

A **hard Jeffrey shift** sets values for p_n regardless of the prior probability P_{n-1}, and so may destroy any information about the partition that was encoded in the prior.

In terms of hard Jeffrey shifts, having the same uncertain evidence at stage n means that $p_n = q_n$.

(M) $P_n(F) = P_{n-1}(F)$ for all $F \in \mathcal{F}_{n-1}$

Theorem (Hutteger). Suppose that $q_n = p_n$, that the sequence Q_n, $n = 1, 2, \ldots$ is uniformly absolutely continuous relative to Q, and that $Q \ll P$. If condition (M) holds, then $d(P_n, Q_n) \to 0$ as $n \to \infty$.

Soft Jeffrey Shifts

$\mathcal{E}_1 = \{E_1, E_2, E_3, E_4\}$ and consider the learning experience given by:

$$(\frac{1}{5} : E_1, \frac{3}{10} : E_2, \frac{1}{2} : E_3, 0 : E_n)$$

Soft Jeffrey Shifts

$\mathcal{E}_1 = \{E_1, E_2, E_3, E_4\}$ and consider the learning experience given by:

$$(\frac{1}{5} : E_1, \frac{3}{10} : E_2, \frac{1}{2} : E_3, 0 : E_n)$$

Consider instead:

- $p_1(E_1) = 2 \cdot P(E_1)$;
- $p_1(E_2) = \frac{1}{2} \cdot P(E_2)$;
- $p_1(E_3) = 5 \cdot P(E_3)$; and
- $p_1(E_4) = 0 \cdot P(E_4)$

If $P(E_1) = \frac{1}{10}$, $P(E_2) = \frac{3}{5}$, $P(E_3) = \frac{1}{10}$ and $P(E_4) = \frac{1}{5}$, then probability kinematics will lead to the same result whether or not it is a hard or “soft” Jeffrey shift.

Do beliefs merge when agents have the same soft uncertain evidence?

Theorem 6.3 in Hutteger shows that this need not be the case.

Do beliefs merge when agents have the same soft uncertain evidence?

Theorem 6.3 in Hutteger shows that this need not be the case.

If Ann and Bob start with different (but mutually absolutely continuous) prior probabilities for infinite sequences of coin flips, and if both observe principle (M) as well as undergo the **same soft Jeffrey shifts**, their posterior degrees of beliefs may not get close to each other in the long run.

“Our results lead to the conclusion that, even under otherwise favorable circumstances, a soft kind of information allows individual rationality to be consistent with sustained disagreement. I don’t think that this is a weakness of the broadly Bayesian approach advocated in this essay. Merging of beliefs happens when it should, i.e., under conditions which may, for example, hold for certain carefully designed scientific investigations. But the claim of merging is not a no-brainer that can be used across the board.”

(Hutteger)

Evidence open to interpretation

C. O'Connor, J. O. Weatherall (2017). *Scientific Polarization*. manuscript.

R. G. Fryer, P. Harms and M. Jackson. *Updating Beliefs When Evidence is Open to Interpretation: Implications for Bias and Polarization*. manuscript, 2017.

Lyme Wars

A decades long scientific debate about:

1. the question of whether Lyme can persist in patients after a short cycle of antibiotics, and
2. the question of whether long term doses of antibiotics are successful in improving the symptoms of Lyme patients.

Lyme Wars

A decades long scientific debate about:

1. the question of whether Lyme can persist in patients after a short cycle of antibiotics, and
2. the question of whether long term doses of antibiotics are successful in improving the symptoms of Lyme patients.

- The values and goals are the same on both sides of the debate.
- They all have access to similar sorts of evidence — they see and treat patients with Lyme and they read the same articles.

Epistemic networks + Jeffry Update

$$P_f(H) = P_i(H \mid E) \cdot P_f(E) + P_i(H \mid \neg E) \cdot P_f(\neg E)$$

Epistemic networks + Jeffry Update

$$P_f(H) = P_i(H \mid E) \cdot P_f(E) + P_i(H \mid \neg E) \cdot P_f(\neg E)$$

1. $P_f(E) = \max(\{1 - d \cdot m \cdot (1 - P_i(E)), 0\})$
2. $P_f(E) = 1 - \min(\{1, d \cdot m \cdot (1 - P_i(E))\})$

    - d is the absolute value of the difference between the agent's credences
    - m is a multiplier that captures how quickly agents become uncertain about the evidence of their peers as their beliefs diverge

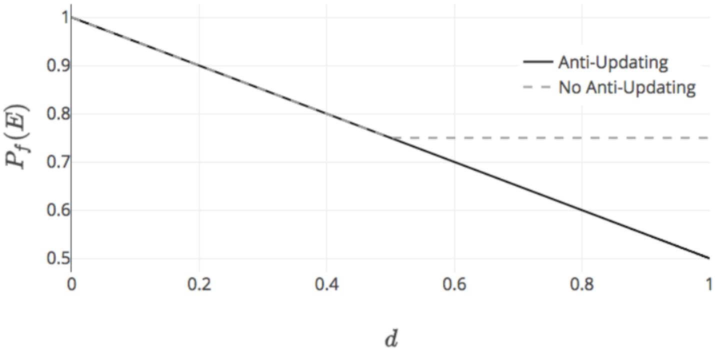

Figure 1: An agent's uncertainty about evidence as a function of distance between credences both for anti-updating and simply ignoring evidence, $m = 2$ and $P(E) = .75$.

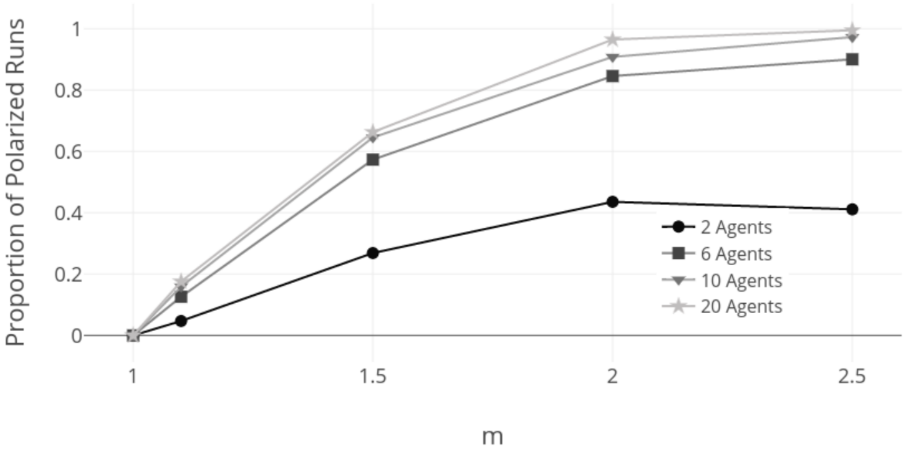

Figure 2: Increasing m increases probability of polarization. $n = 50$, $p_B = .7$.

Participants presented with one of two questions:

- Do you think the death penalty deters (stops) people from committing murder?
- Do you think human activity is the cause of rising temperatures?

Respondents were asked to answer on a scale from -8 (I am certain that the death penalty does NOT deter people from committing murder/I am certain that human activity is NOT the cause of increasing temperatures) to 8 (I am certain that the death penalty DOES deter people from committing murder/I am certain that human activity IS the cause of increasing temperatures).

Next, participants were presented short summaries. After each summary, they were asked whether they thought the summary provided evidence for or against the topic.

The summaries came from articles that provided a variety of conclusions and would be easy to understand. After selecting articles, they redacted the abstracts and introductions into a short summary of each article (keeping each summary at more or less the same length and level of readability). An initial survey labeled the summaries as PRO, CON or UNCLEAR.

Finally the participants were asked the original question again.

- State: $\{A, B\}$
- Signals: $\{a, b, ab, \emptyset\}$
    - a is a signal for state A
    - b is a signal for state B
    - ab is an ambiguous signal
    - $\emptyset$ represents the lack of a signal
- λ_0: prior probability
- q: probability of observing a signal
- p: probability that an agent receives signal a if the state is A and b if the state is B
- π: probability that a signal "becomes" ambiguous.

Theorem. A Bayesian-updating agent who forms beliefs conditional upon the full sequence of signals has a posterior that converges to place probability 1 on the correct state, almost surely.

Proposition. Suppose that a nontrivial fraction of experiences are open to interpretation so that $\pi > \frac{p-1/2}{p}$. Consider two interpretative agents 1 and 2 who both use the **maximum likelihood rule** but have differing priors: agent 1's prior is that A is more likely (so 1 has a prior $\lambda_0 > 1/2$) and agent 2's prior is that B is more likely (so 2 has a prior $\lambda_0 < 1/2$).

Proposition. Suppose that a nontrivial fraction of experiences are open to interpretation so that $\pi > \frac{p-1/2}{p}$. Consider two interpretative agents 1 and 2 who both use the **maximum likelihood rule** but have differing priors: agent 1's prior is that A is more likely (so 1 has a prior $\lambda_0 > 1/2$) and agent 2's prior is that B is more likely (so 2 has a prior $\lambda_0 < 1/2$).

Let the two agents see exactly the same sequence of signals. With a positive probability that tends to 1 in π the two agents will end up polarized with 1's posterior tending to 1 and 2's posterior tending to 0. With a positive probability tending to 0 in π the two agents will end up with the same (possibly incorrect) posterior tending to either 0 or 1

Theorem. (Bayesian Consensus-or-Polarization Law). Let P and Q be any two probabilities and suppose that P is **absolutely continuous with respect to Q to degree δ**. Let M be the event that P and Q merge, and let L be the event that P and Q polarize in the limit. If P shares an increasing and complete sequence of evidence with Q, then $P(M) = \delta$ and $P(L) = 1 - \delta$.

M. Nielsen and R. Stewart. *Persistent Disagreement and Polarization in a Bayesian Setting*. manuscript, 2018.

Concluding Remarks

- ✓ **Monday** Representing judgements; Introduction to judgement aggregation; Aggregation paradoxes I
- ✓ **Tuesday** Aggregation paradoxes II, Axiomatic characterizations of aggregation methods I
- ✓ **Wednesday** Axiomatic characterizations of probabilistic opinions
- ✓ **Thursday** Pooling imprecise probabilities; Distance-based characterizations; Aumann's agreeing to disagree theorem and related results
- ✓ **Friday** Merging of probabilistic opinions (Blackwell-Dubins Theorem); Belief polarization; Concluding remarks

Some topics we missed

- Aggregating incoherent probabilities

R. Pettigrew. *Aggregating incoherent agents who disagree*. Synthese, 2017.

- Aggregating causal models

D. Alrajeh, H. Chockler, and J. Halpern. *Combining experts' causal judgments*. Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18), 2017.

- Diversity trumps ability theorem (Hong-Page Theorem)

L. Hong and S. E. Page. *Groups of diverse problem solvers can outperform groups of high-ability problem solvers*. PNAS, 101(46), pp. 16385-16389, 2004.

Thank You!

pacuit.org

epacuit@umd.edu