Agreeing to Disagree and Dilation

Jiji Zhang, Hailin Liu and Teddy Seidenfeld
jijizhang@ln.edu.hk liuhlin3@mail.sysu.edu.cn teddy@stat.cmu.edu

Lingnan University 嶺南大學 · 中山大學 SUN YAT-SEN UNIVERSITY · Carnegie Mellon University

1. Introduction & Take-home Message

Aumann (1976) gives sufficient conditions when two (precise) Bayesian agents with the same prior over a measurable space cannot agree to disagree on posteriors of a hypothesis H: Their posteriors for H must be equal if they are commonly known, along with their personal information sets leading to these posteriors.

- Geanakoplos and Polemarchakis (1982) extend Aumman's result to a setting where the agents make their credences common knowledge by communication.
- Kajii and Ui (2005) and Carvajal and Correia-da-Silva (2010) generalize Aumman's result to a setting of imprecise priors.
- **We investigate Aumann's result in a combination of these two more general settings, and show that the interesting and anomalous phenomenon known as *dilation* is the key obstacle to reaching agreements via communicating posteriors by Bayesian agents with imprecise priors.**

2. Aumann's Agreement Theorem

- Suppose two agents have the same (precise) prior, p, over a measurable space $(\Omega, \mathcal{A})$.
- Agent i learns (privately) the value of a partition of Ω, $\mathcal{P}_i$, and updates by Bayesian conditioning.
- All these are *commonly* known: each agent knows them, knows that each knows them, knows that each knows that each knows them, ... *ad infinitum*.
- Then, if it is common knowledge that agent 1's posterior of an event H is p_1 and agent 2's posterior of H is p_2, then $p_1 = p_2$. That is, the agents cannot *agree to disagree*!

Example of Agreement

$\mathcal{P}^1$ \ $\mathcal{P}^2$		
	w_1	w_2
	w_3	w_4

$\mathcal{P}^1 = \{\{w_1, w_2\}, \{w_3, w_4\}\}$
$\mathcal{P}^2 = \{\{w_1, w_3\}, \{w_2, w_4\}\}$
$H = \{w_1, w_4\}$

- Suppose $p = (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$, and suppose agent 1 learns $\{w_1, w_2\}$ while agent 2 learns $\{w_1, w_3\}$.
- So, $p_1(H) = p(H|\{w_1, w_2\}) = \frac{1}{2}$. Similarly, $p_2(H) = p(H|\{w_1, w_3\}) = \frac{1}{2}$.
- Note: it is common knowledge that $p_1(H) = \frac{1}{2}$ and that $p_2(H) = \frac{1}{2}$.

Example of Disagreement

$\mathcal{P}^1$ \ $\mathcal{P}^2$		
	w_1	w_2
	w_3	w_4

$\mathcal{P}^1 = \{\{w_1, w_2\}, \{w_3, w_4\}\}$
$\mathcal{P}^2 = \{\{w_1, w_3\}, \{w_2, w_4\}\}$
$H' = \{w_4\}$

- Suppose $p = (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$, and suppose agent 1 learns $\{w_1, w_2\}$ while agent 2 learns $\{w_2, w_4\}$.
- Then, $p_1(H') = p(H'|\{w_1, w_2\}) = 0 \neq \frac{1}{2} = p(H'|\{w_2, w_4\}) = p_2(H')$.
- This is *not* agreeing to disagree, for they will reach an agreement if they share posteriors.

3. Example of Disagreement with IP

$\mathcal{P}^1$		
	w_1	w_2
	w_3	w_4

$\mathcal{P}^1 = \{\{w_1, w_2\}, \{w_3, w_4\}\}$
$\mathcal{P}^2 = \{\Omega\}$; $H = \{w_1, w_4\}$

- Common imprecise prior (an ϵ-contaminated class): $\mathbf{Q} = \{0.8p + 0.2q \mid q \in \Lambda\}$, where Λ is the set of all distributions over $\mathcal{A}$.
- $\mathbf{Q}_1(H) = \mathbf{Q}(H|\{w_1, w_2\}) = [\frac{1}{3}, \frac{2}{3}]$, whereas $\mathbf{Q}_2(H) = \mathbf{Q}(H|\Omega) = [\frac{2}{5}, \frac{3}{5}]$.
- In this case $\mathbf{Q}_1(H)$ and $\mathbf{Q}_2(H)$ are common knowledge: the agents agree to disagree!

Dilation of IP

- Observe that H is *dilated* by $\mathcal{P}^1$, in the sense that for every $E \in \mathcal{P}^1$, $[\underline{\mathbf{Q}}(H|E), \overline{\mathbf{Q}}(H|E)] = [\frac{1}{3}, \frac{2}{3}]$ strictly contains $[\underline{\mathbf{Q}}(H), \overline{\mathbf{Q}}(H)] = [\frac{2}{5}, \frac{3}{5}]$.
- In general, a set of mutually disjoint events $\mathcal{E}$ *dilates* an event A with respect to $\mathbf{Q}$ if for every $E \in \mathcal{E}$, the interval $[\underline{\mathbf{Q}}(A|E), \overline{\mathbf{Q}}(A|E)]$ strictly contains the interval $[\underline{\mathbf{Q}}(A|\bigcup\mathcal{E}), \overline{\mathbf{Q}}(A|\bigcup\mathcal{E})]$.
- It is no accident that dilation occurs when the agents can agree to disagree!

4. A Procedure of Communicating Imprecise Posteriors

Suppose the true state is w. The initial common knowledge is $\mathcal{C}_0 = \mathcal{P}(w)$, where $\mathcal{P}$ is the finest common coarsening of $\mathcal{P}^1$ and $\mathcal{P}^2$.

- Step 0: Agent i updates her credence of H to $\mathbf{Q}_0^i(H) = \mathbf{Q}(H|\mathcal{P}^i(w))$. Let $\mathcal{P}_0^i = \{E \in \mathcal{P}^i \mid E \cap \mathcal{C}_0 \neq \emptyset\}$.
- Step $n+1$: They announce $\mathbf{Q}_n^1(H)$ and $\mathbf{Q}_n^2(H)$, respectively. Let

$$\mathcal{N}_{n+1}^i = \{E \in \mathcal{P}_n^i \mid \mathbf{Q}(H|E \cap \mathcal{C}_n) = \mathbf{Q}_n^i(H)\}$$
$$\mathcal{C}_{n+1} = \bigcup \mathcal{N}_{n+1}^1 \cap \bigcup \mathcal{N}_{n+1}^2, \ \& \ \mathcal{P}_{n+1}^i = \{E \in \mathcal{N}_{n+1}^i \mid E \cap \mathcal{C}_{n+1} \neq \emptyset\}.$$

If $\mathcal{P}_{n+1}^i = \mathcal{P}_n^i$ (or $\mathcal{C}_{n+1} = \mathcal{C}_n$), neither agent learns new information and the procedure stops; otherwise, agent i updates credence of H to $\mathbf{Q}_{n+1}^i(H) = \mathbf{Q}(H|\mathcal{P}^i(w) \cap \mathcal{C}_{n+1})$, and enters the next step.

5. A Generalization of Aumann's Agreement Theorem

- In the absence of dilation, two agents are guaranteed to reach consensus on lower and upper probabilities by communicating posteriors. More formally:

Theorem: Suppose the above procedure stops at step $m+1$. If for both $i = 1, 2$, $\{E \cap \mathcal{C}_m \mid E \in \mathcal{P}_m^i\}$ does not dilate H, then $\underline{\mathbf{Q}}(H|\mathcal{P}^1(w) \cap \mathcal{C}_m) = \underline{\mathbf{Q}}(H|\mathcal{P}^2(w) \cap \mathcal{C}_m)$ and $\overline{\mathbf{Q}}(H|\mathcal{P}^1(w) \cap \mathcal{C}_m) = \overline{\mathbf{Q}}(H|\mathcal{P}^2(w) \cap \mathcal{C}_m)$.

- It is easy to show that the above result still holds if at each step the agents communicate only lower and upper probabilities, instead of the whole sets of probablities.

A Result on Full Agreement

- Under some common assumptions, lower and upper probabilities are sufficient to identify the full set. For example,

Corollary 1: Suppose $\mathbf{Q}$ is closed and connected and the procedure of communicating posteriors stops at step $m+1$. If for both $i = 1, 2$, $\{E \cap \mathcal{C}_m \mid E \in \mathcal{P}_m^i\}$ does not dilate H, then $\mathbf{Q}(H|\mathcal{P}^1(w) \cap \mathcal{C}_m) = \mathbf{Q}(H|\mathcal{P}^2(w) \cap \mathcal{C}_m)$.

A Corollary for Density Ratio Priors

- Seidenfeld and Wasserman (1993, Theorem 4.1) showed that the density ratio priors are dilation-immune. Thus we have:

Corollary 2: Suppose two agents start with a common density ratio prior $\mathbf{Q}$ and carry out the procedure of communicating posteriors. Suppose the procedure stops at step $m+1$. Then $\mathbf{Q}(H|\mathcal{P}^1(w) \cap \mathcal{C}_m) = \mathbf{Q}(H|\mathcal{P}^2(w) \cap \mathcal{C}_m)$.

Partial Agreement with IP

- If we just consider partial agreement in the sense of a non-empty intersection of sets of posteriors, we can drop the assumption of connectedness.

Corollary 3: Suppose $\mathbf{Q}$ is closed and the procedure of communicating posteriors stops at step $m+1$. If for both $i = 1, 2$, $\{E \cap \mathcal{C}_m \mid E \in \mathcal{P}_m^i\}$ does not dilate H, then $\mathbf{Q}(H|\mathcal{P}^1(w) \cap \mathcal{C}_m) \cap \mathbf{Q}(H|\mathcal{P}^2(w) \cap \mathcal{C}_m) \neq \emptyset$.

6. Concluding Remarks and Further Questions

- The presence of dilation is necessary for agreeing to disagree on the lower or upper posterior.
- To put it differently, *dilation-averse* agents cannot agree to disagree on the lower or upper posterior.

- What about agents whose priors agree only *partially*?
- What about other updating rules, e.g., the Dempster-Shafer rule?