Aggregating Judgements: Logical and Probabilistic Approaches

Lecture 4

Eric Pacuit
Department of Philosophy
University of Maryland
pacuit.org

August 9, 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; Merging of probabilistic opinions (Blackwell-Dubins Theorem); Aumann's agreeing to disagree theorem and related results
- **Friday** Belief polarization; Diversity trumps ability theorem (The Hong-Page Theorem)

Aggregating imprecise probabilities

Imprecise Probabilities

1. What is the probability that a fair coin will land hands?

Imprecise Probabilities

1. What is the probability that a fair coin will land hands?
2. What is the probability of a coin of unknown bias will land heads?

Imprecise Probabilities

1. What is the probability that a fair coin will land hands?
2. What is the probability of a coin of unknown bias will land heads?

Ellsberg Paradox

Ellsberg Paradox

	30	60	
Lotteries	Blue	Yellow	Green
L_1	$1M$	0	0
L_2	0	$1M$	0
L_3	$1M$	0	$1M$
L_4	0	$1M$	$1M$

Ellsberg Paradox

	30	60	
Lotteries	Blue	Yellow	Green
L_1	$1M$	0	0
L_2	0	$1M$	0
L_3	$1M$	0	$1M$
L_4	0	$1M$	$1M$

$$L_1 \succeq L_2 \text{ iff } L_3 \succeq L_4$$

Indeterminate Probability

- Allow probability functions to take on sets of values instead of a single value
- Work with sets of probabilities rather than a single probability

Precisification Given a function $\sigma : \Sigma \to \wp([0, 1])$, a probability function $p : \Sigma \to [0, 1]$ of σ if and only if $p(A) \in \sigma(A)$ for each $A \in \Sigma$.

Indeterminate Probability A function $\sigma : \Sigma \to \wp([0, 1])$ such that whenever $x \in \sigma(A)$ there is some precisifcation of σ, p for which $p(A) = x$.

Convexity A class of probability functions Π is **convex** if and only if whenever $p, q \in \Pi$, every mixture of p and q is in Π as well. I.e., $\alpha p + (1 - \alpha)q \in \Pi$ for all $\alpha \in (0, 1)$.

Proposition. If P is convex with σ it ambiguation, then $\sigma(A)$ is an interval for each A.

IP Pooling

$$F : \mathcal{P}^n \to \wp(\mathcal{P})$$

R. T. Stewart and I. Ojea Quintana. *Probabilistic Opinion Pooling with Imprecise Probabilities*. Journal of Philosophical Logic, 47(1), pp. 17 - 45, 2018.

IP Pooling

$$F : \mathcal{P}^n \to \wp(\mathcal{P})$$

R. T. Stewart and I. Ojea Quintana. *Probabilistic Opinion Pooling with Imprecise Probabilities*. Journal of Philosophical Logic, 47(1), pp. 17 - 45, 2018.

IP Pooling: For each $\mathbf{p} = (p_1, \ldots, p_n)$, $\quad F(\mathbf{p}) = conv\{p_i \mid i = 1, \ldots, n\}$,

where $conv(X)$ is the **convex hull** of a set X of probabilities.

Proposition (Stewart and Ojea Quintana) Convex IP pooling functions satisfy event-wise independence, unanimity preservation (and other properties of linear pooling studied in the literature)

Proposition (Stewart and Ojea Quintana) Convex IP pooling functions satisfy event-wise independence, unanimity preservation (and other properties of linear pooling studied in the literature)

Proposition (Stewart and Ojea Quintana) Convex IP pooling functions are externally Bayesian.

Proposition (Stewart and Ojea Quintana) Convex IP pooling functions satisfy event-wise independence, unanimity preservation (and other properties of linear pooling studied in the literature)

Proposition (Stewart and Ojea Quintana) Convex IP pooling functions are externally Bayesian.

Proposition (Stewart and Ojea Quintana) Convex IP pooling functions are not individualwise Bayesian.

Individualwise Bayesian: For all $\mathbf{p} = (p_1, \ldots, p_n)$ and likelihood functions L, $F^L(\mathbf{p}) = F(p_1, \ldots, p_k^L, \ldots, p_n)$.

Aggregating IP

S. Moral and J. Del Sagrado. *Aggregation of imprecise probabilities*. In Aggregation and Fusion of Imperfect Information, pp. 162 - 188. Springer, 1998.

R. F. Nau. *The aggregation of imprecise probabilities*. Journal of Statistical Planning and Inference 105 (1), pp. 265 - 282, 2002.

(Among others...)

Distance-based characterization of aggregation methods

M. Miller and D. Osherson. *Methods for distance-based judgment aggregation*. Social Choice and Welfare, 32(4), pp. 575 - 601, 2009.

G. Pigozzi. *Belief merging and the discursive dilemma: an argument-based account to paradoxes of judgment aggregation*. Synthese, 152(2), pp. 285?298, 2006.

J. Lang, G. Pigozzi, M. Slavkovik, and L. van der Torre. *Judgment aggregation rules based on minimization*. In Proceedings of TARK, pp. 238 - 246, 2011.

Independence?

Independence: For any $p \in A$ and all $\mathbf{J} = (J_1, \ldots, J_n)$ and $\mathbf{J}^* = (J_1^*, \ldots, J_n^*)$ in the domain of F,

$$\text{if [for all } i \in N,\ p \in J_i \text{ iff } p \in J_i^*]$$
$$\text{then } [p \in F(\mathbf{J}) \text{ iff } p \in F(\mathbf{J}^*)\ ].$$

Finding a group judgement set that is *as close as possible* to the group judgements will not satisfy independence.

Given $(J_1, \ldots, J_n)$, select the set consistent and complete J that minimizes the total distance from the individual judgement sets: find J such that $\sum_{i \in N} d(J, J_i)$ is minimized, where $d(J, J_i)$ is the *distance* between J and J_i

Given $(J_1, \ldots, J_n)$, select the set consistent and complete J that minimizes the total distance from the individual judgement sets: find J such that $\sum_{i \in N} d(J, J_i)$ is minimized, where $d(J, J_i)$ is the *distance* between J and J_i

Hamming Metric: $d(J, J') =$ the number of propositions for which J and J' disagree

	p	q	$p \wedge q$
1	T	T	T
2	T	F	F
3	F	T	F
Majority	T	T	F

	p	q	$p \wedge q$
1	T	T	T
2	T	F	F
3	F	T	F
Majority	T	T	F
Premise	T	T	T

	p	q	$p \wedge q$
1	T	T	T
2	T	F	F
3	F	T	F
Majority	T	T	F
Premise	T	T	T
Hamming 1	F	T	F
Hamming 2	T	F	F

Differing on $\{a, b \wedge c\}$ may be considered more consequential than differing on $\{a, a \wedge b\}$.

Differing on $\{a, b \wedge c\}$ may be considered more consequential than differing on $\{a, a \wedge b\}$.

Let $\mathcal{F}$ be the set of *all* judgement sets and $\mathcal{F}^{\circ}$ the set of all consistent judgement sets.

$d : \mathcal{F} \times \mathcal{F} \rightarrow \mathbb{R}$

Axiom 1 $d(A, B) = 0$ iff $A = B$
Axiom 2 $d(A, B) = d(B, A)$
Axiom 3 $d(A, B) \leq d(A, C) + d(C, B)$

$$d_H(\{p, q, p \wedge q\}, \{p, \neg q, \neg(p \wedge q)\}) = 2$$

$d_H(\{p, q, p \wedge q\}, \{p, \neg q, \neg(p \wedge q)\}) = 2$

Shouldn't $d(\{p, q, p \wedge q\}, \{p, \neg q, \neg(p \wedge q)\}) = 1$?

C. Duddy and A. Piggins. *A measure of distance between judgement sets*. Social Choice and Welfare, 39, pp. 855 - 867, 2012.

Duddy and Piggins Measure

Judgement set C is between judgement sets A and B if A, B and C are distinct and, on each proposition C agrees with A or with B (or both). (*C is a compromise between A and B*)

Duddy and Piggins Measure

Judgement set C is between judgement sets A and B if A, B and C are distinct and, on each proposition C agrees with A or with B (or both). (*C is a compromise between A and B*)

Draw a graph where the nodes are possible judgement sets and there is an edge between A and B provided there is no judgement set between them.

The distance between A and B is the length of the shortest path from A to B.

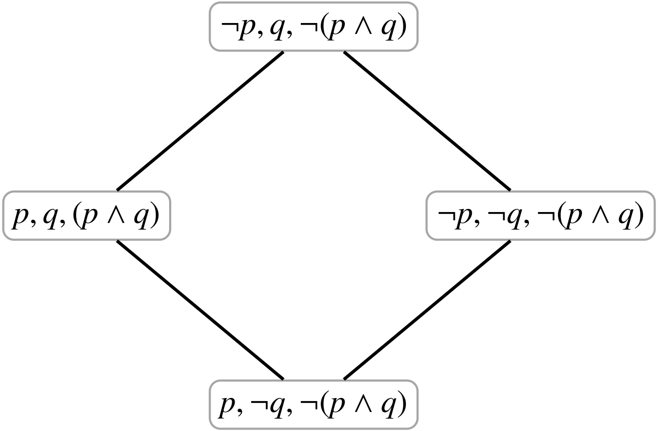
$\neg p, q, \neg(p \wedge q)$
$p, q, (p \wedge q)$
$\neg p, \neg q, \neg(p \wedge q)$
$p, \neg q, \neg(p \wedge q)$

Axioms

Axiom 1 $d(A, B) = 0$ iff $A = B$
Axiom 2 $d(A, B) = d(B, A)$
Axiom 3 $d(A, B) \leq d(A, C) + d(C, B)$

Axioms

Axiom 1 $d(A, B) = 0$ iff $A = B$
Axiom 2 $d(A, B) = d(B, A)$
Axiom 3 $d(A, B) \leq d(A, C) + d(C, B)$

For all A, B, C, C is between A and B provided $A \neq B \neq C$ and $(A \cap B) \subset C$.

Axioms

Axiom 1 $d(A, B) = 0$ iff $A = B$
Axiom 2 $d(A, B) = d(B, A)$
Axiom 3 $d(A, B) \leq d(A, C) + d(C, B)$

For all A, B, C, C is between A and B provided $A \neq B \neq C$ and $(A \cap B) \subset C$.

Axiom 4 If there is a judgement set between A and B then there exists C different from A and B such that $d(A, B) = d(A, C) + d(C, B)$

Axioms

Axiom 1 $d(A, B) = 0$ iff $A = B$
Axiom 2 $d(A, B) = d(B, A)$
Axiom 3 $d(A, B) \leq d(A, C) + d(C, B)$

For all A, B, C, C is between A and B provided $A \neq B \neq C$ and $(A \cap B) \subset C$.

Axiom 4 If there is a judgement set between A and B then there exists C different from A and B such that $d(A, B) = d(A, C) + d(C, B)$

Axiom 5 If there is no judgement set between A and B with $A \neq B$ then $d(A, B) = 1$

Theorem (Duddy & Piggins) The previously defined metric is the unique metric satisfying Axioms 1 - 5.

	p	q	$p \wedge q$
1	T	T	T
2	T	F	F
3	F	T	F
Majority	T	T	F
Premise	T	T	T
Hamming 1	F	T	F
Hamming 2	T	F	F

	p	q	$p \wedge q$
1	T	T	T
2	T	F	F
3	F	T	F
Majority	T	T	F
Premise	T	T	T
Hamming 1	F	T	F
Hamming 2	T	F	F
DP-metric	T	T	T

Let J be a profile.

Find profiles J^* such that $\sum_i d(J_i, J)$ is minimized

vs.

Find profiles J^* that minimizes $\sum d(J, J^*)$

where profiles J and J', $d(J, J') = \sum_{i \leq n} d(J_i, J'_i)$

M. Miller and D. Osherson. *Methods for distance-based judgement aggregation*. Social Choice and Welfare, 32, pgs. 575 - 601, 2009.

For a profile P, $M(P) \in \mathcal{F}$ the judgement set resulting from majority rule. P is majority consistent provided $M(P) \in \mathcal{F}^\circ$

Fix a metric d and a profile $J \in \mathcal{F}^\circ$

For a profile P, $M(P) \in \mathcal{F}$ the judgement set resulting from majority rule. P is majority consistent provided $M(P) \in \mathcal{F}^\circ$

Fix a metric d and a profile $J \in \mathcal{F}^\circ$

- $Full_d(J)$ is the collection of $M(J') \in \mathcal{F}^\circ$ such that J' minimizes $d(J, J')$ over all majority consistent profiles J' in $\mathcal{F}^\circ$

For a profile P, $M(P) \in \mathcal{F}$ the judgement set resulting from majority rule. P is majority consistent provided $M(P) \in \mathcal{F}^\circ$

Fix a metric d and a profile $J \in \mathcal{F}^\circ$

- $Full_d(J)$ is the collection of $M(J') \in \mathcal{F}^\circ$ such that J' minimizes $d(J, J')$ over all majority consistent profiles J' in $\mathcal{F}^\circ$
- $Output_d(J)$ is the collection of $M(J') \in \mathcal{F}^\circ$ such that J' minimizes $d(J, J')$ over all majority profiles J' in $\mathcal{F}$ (*allowing inconsistencies*)

For a profile P, $M(P) \in \mathcal{F}$ the judgement set resulting from majority rule. P is majority consistent provided $M(P) \in \mathcal{F}^\circ$

Fix a metric d and a profile $J \in \mathcal{F}^\circ$

- $Full_d(J)$ is the collection of $M(J') \in \mathcal{F}^\circ$ such that J' minimizes $d(J, J')$ over all majority consistent profiles J' in $\mathcal{F}^\circ$
- $Output_d(J)$ is the collection of $M(J') \in \mathcal{F}^\circ$ such that J' minimizes $d(J, J')$ over all majority profiles J' in $\mathcal{F}$ (*allowing inconsistencies*)
- $Endpoint_d(J)$ is the collection of $K \in \mathcal{F}^\circ$ that minimize $d(J, K)$ over all majority consistent profiles J'

For a profile P, $M(P) \in \mathcal{F}$ the judgement set resulting from majority rule. P is majority consistent provided $M(P) \in \mathcal{F}^\circ$

Fix a metric d and a profile $J \in \mathcal{F}^\circ$

- $Full_d(J)$ is the collection of $M(J') \in \mathcal{F}^\circ$ such that J' minimizes $d(J, J')$ over all majority consistent profiles J' in $\mathcal{F}^\circ$
- $Output_d(J)$ is the collection of $M(J') \in \mathcal{F}^\circ$ such that J' minimizes $d(J, J')$ over all majority profiles J' in $\mathcal{F}$ (*allowing inconsistencies*)
- $Endpoint_d(J)$ is the collection of $K \in \mathcal{F}^\circ$ that minimize $d(J, K)$ over all majority consistent profiles J'
- $Prototype_d(J)$ is the collection of $K \in \mathcal{F}^\circ$ that minimize $\sum_{i \leq n} d(J_i, K)$ over all $K \in \mathcal{F}^\circ$

For J, K let $Ham(J, K)$ denote the Hamming distance (the number of items on which J and K disagree)

$$d(J, K) = \begin{cases} 0.9 & \text{if } J \text{ and } K \text{ disagree only on } a \wedge b \\ \sqrt{Ham(p, q)} & \text{otherwise} \end{cases}$$

	a	b	$a \wedge b$	a	b	$a \wedge b$	a	b	$a \wedge b$
1	T	T	T	T	T	T	T	T	T
2	T	T	T	T	T	T	T	T	T
3	T	F	F	T	F	F	**T**	**F**	**T**
4	T	F	F	T	F	F	T	F	F
5	F	T	F	**F**	**F**	**F**	F	T	F
M	T	T	F	T	F	F	T	T	T

	a	b	$a \wedge b$	a	b	$a \wedge b$	a	b	$a \wedge b$
1	T	T	T	T	T	T	T	T	T
2	T	T	T	T	T	T	T	T	T
3	T	F	F	T	F	F	**T**	**F**	**T**
4	T	F	F	T	F	F	T	F	F
5	F	T	F	**F**	**F**	**F**	F	T	F
M	T	T	F	T	F	F	T	T	T

- $Full_d(J) = TFF$ $(d(FTF, FFF) = 1)$

	a	b	$a \wedge b$	a	b	$a \wedge b$	a	b	$a \wedge b$
1	T	T	T	T	T	T	T	T	T
2	T	T	T	T	T	T	T	T	T
3	T	F	F	T	F	F	**T**	**F**	**T**
4	T	F	F	T	F	F	T	F	F
5	F	T	F	**F**	**F**	**F**	F	T	F
M	T	T	F	T	F	F	T	T	T

- $Full_d(J) = TFF \ (d(FTF, FFF) = 1)$
- $Output_d(J) = TTT \ (d(TFF, TFT) = 0.9)$

	a	b	$a \wedge b$	a	b	$a \wedge b$	a	b	$a \wedge b$
1	T	T	T	T	T	T	T	T	T
2	T	T	T	T	T	T	T	T	T
3	T	F	F	T	F	F	**T**	**F**	**T**
4	T	F	F	T	F	F	T	F	F
5	F	T	F	**F**	**F**	**F**	F	T	F
M	T	T	F	T	F	F	T	T	T

- $Full_d(J) = TFF$ $(d(FTF, FFF) = 1)$
- $Output_d(J) = TTT$ $(d(TFF, TFT) = 0.9)$
- $Endpoint_d(J) = TTT$ $(d(TTF, TTT) = 0.9)$

	a	b	$a \wedge b$	a	b	$a \wedge b$	a	b	$a \wedge b$
1	T	T	T	T	T	T	T	T	T
2	T	T	T	T	T	T	T	T	T
3	T	F	F	T	F	F	**T**	**F**	**T**
4	T	F	F	T	F	F	T	F	F
5	F	T	F	**F**	**F**	**F**	F	T	F
M	T	T	F	T	F	F	T	T	T

- $Full_d(J) = TFF$ $(d(FTF, FFF) = 1)$
- $Output_d(J) = TTT$ $(d(TFF, TFT) = 0.9)$
- $Endpoint_d(J) = TTT$ $(d(TTF, TTT) = 0.9)$
- $Prototype_d(J) = \{TTT, TFF\}$ $(\sum_i d(J_i, TTT) = 3\sqrt{2}$, $\sum_i d(J_i, TFF) = 3\sqrt{2}$, $\sum_i d(J_i, FTF) = 4\sqrt{2}$, $\sum_i d(J_i, FFF) = 2\sqrt{3} + 3)$

Rational Disagreement

Starting with the same premises, using (for example) first-order logic, two agents cannot disagree about whether a conclusion follows.

Starting with the same probability, using (for example) strict conditionalization, two agents cannot disagree about their posterior probability given the same evidence.

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)

Aumann's Agreeing to Disagree Theorem. Suppose that n agents share a common prior and have different private information. If there is common knowledge of the posteriors of a fixed event, then the posteriors must be equal.

Robert Aumann. *Agreeing to Disagree*. Annals of Statistics **4(6)**, pgs. 1236-1239 (1976).

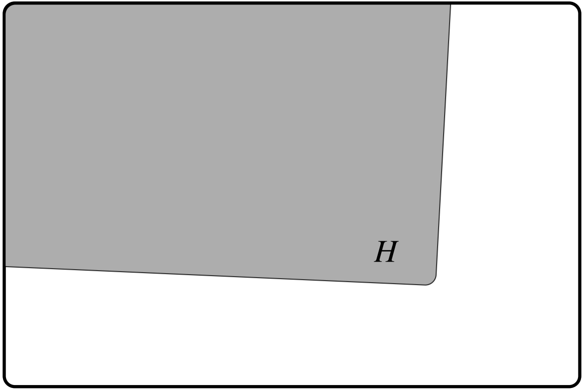

An **event/proposition** is a (definable) subset $H \subseteq W$.

A σ**-algebra** is the collection of events/propositions
(closed under countable unions and complementation)

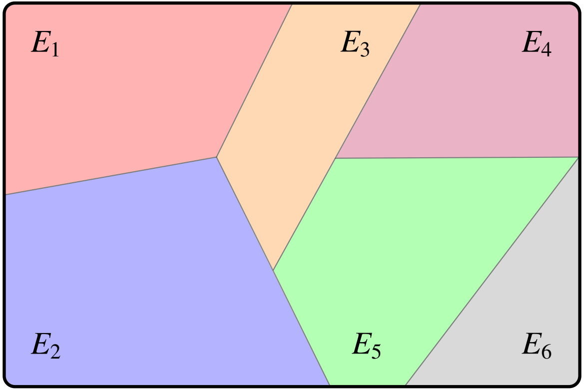

An **experiment/question/set of signals** is a partition $\mathcal{E}$ on W.

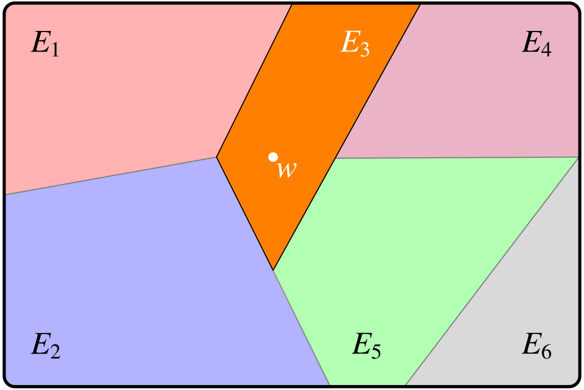

An **experiment/question/set of signals** is a partition $\mathcal{E}$ on W.

If $w \in W$, let $\mathcal{E}[w] = E$ where $w \in E \in \mathcal{E}$.

E.g, if $\mathcal{E} = \{E_1, E_2, E_3, E_4, E_5, E_6\}$, then $\mathcal{E}[w] = E_3$

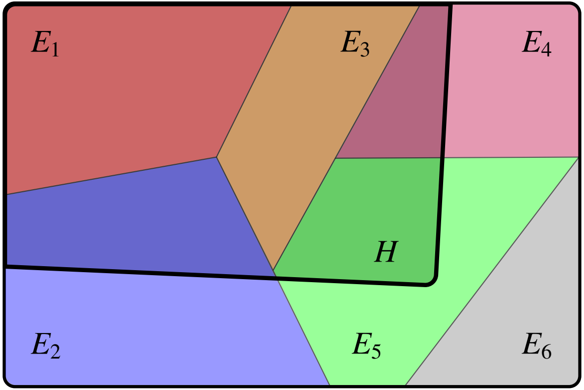

$K_{\mathcal{E}} : \wp(W) \to \wp(W)$, where for $H \subseteq W$,

$$K_{\mathcal{E}}(H) = \{w \mid \mathcal{E}[w] \subseteq H\}$$

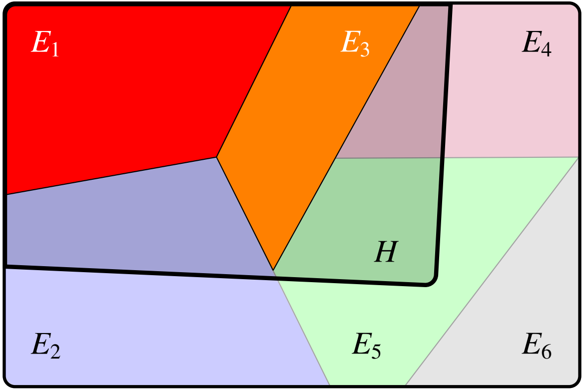

$$K_{\mathcal{E}}(H) \quad = \quad E_1 \cup E_3$$

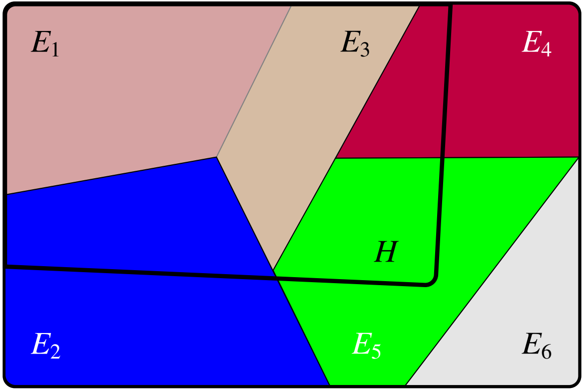

$$
\begin{aligned}
K_{\mathcal{E}}(H) &= E_1 \cup E_3 \\
-K_{\mathcal{E}}(H) \cap -K_{\mathcal{E}}(-H) &= E_2 \cup E_4 \cup E_5
\end{aligned}
$$

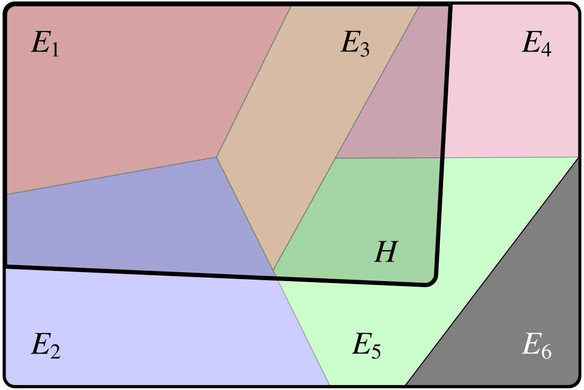

$$
\begin{aligned}
K_{\mathcal{E}}(H) &= E_1 \cup E_3 \\
-K_{\mathcal{E}}(H) \cap -K_{\mathcal{E}}(-H) &= E_2 \cup E_4 \cup E_5 \\
K_{\mathcal{E}}(-H) &= E_6
\end{aligned}
$$

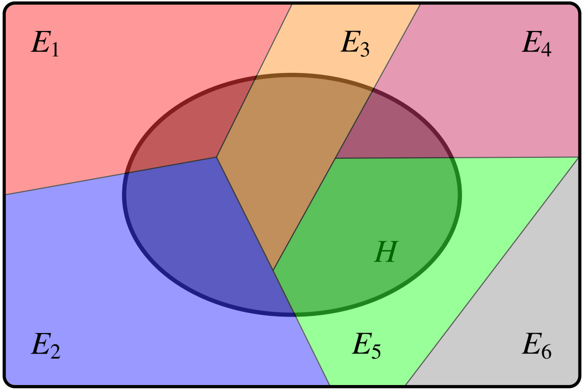

If P is a probability on W (with respect to a σ-algebra $\mathcal{F}$)

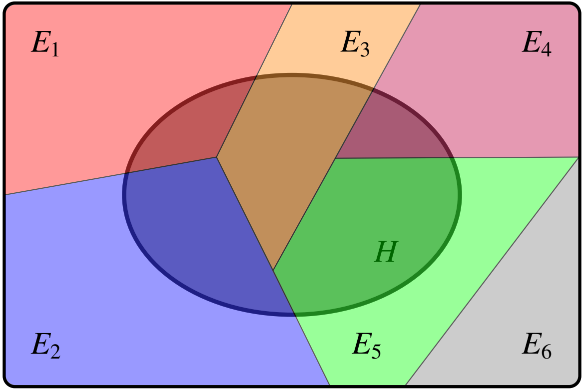

If P is a probability on W (with respect to a σ-algebra $\mathcal{F}$)

The posterior at w with respect to $\mathcal{E}$ is $P_{\mathcal{E},w}(H) = P(H \mid \mathcal{E}[w])$

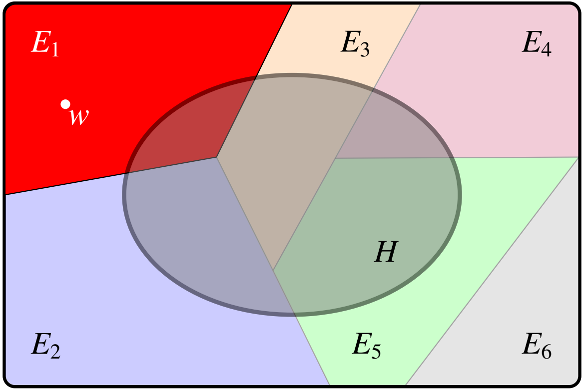

If P is a probability on W (with respect to a σ-algebra $\mathcal{F}$)

The posterior at w with respect to $\mathcal{E}$ is $P_{\mathcal{E},w}(H) = P(H \mid \mathcal{E}[w])$

E.g., $P_{\mathcal{E},w}(H) = P(H \mid E_1)$

A basic result about probabilities.

For any finite partition $\mathcal{E} = \{E_i\}$ of W and an event H,

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

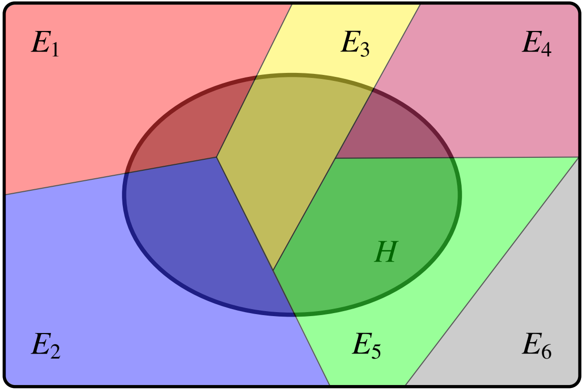

$$P(H) \quad = \quad P(H \cap E_1) + \cdots + P(H \cap E_6)$$

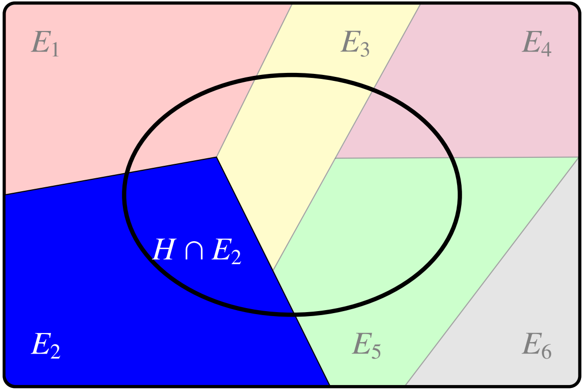

$$P(H) \quad = \quad P(H \cap E_1) + P(H \cap E_2) + \cdots + P(H \cap E_6)$$

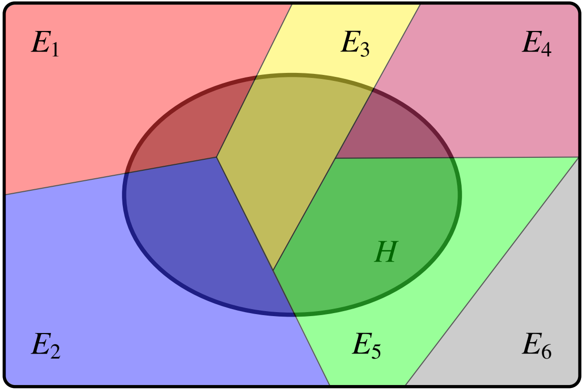

$$
\begin{aligned}
P(H) &= P(H \cap E_1) + \cdots + P(H \cap E_6) \\
&= \tfrac{P(E_1)}{P(E_1)} P(H \cap E_1) + \cdots + \tfrac{P(E_6)}{P(E_6)} P(H \cap E_6)
\end{aligned}
$$

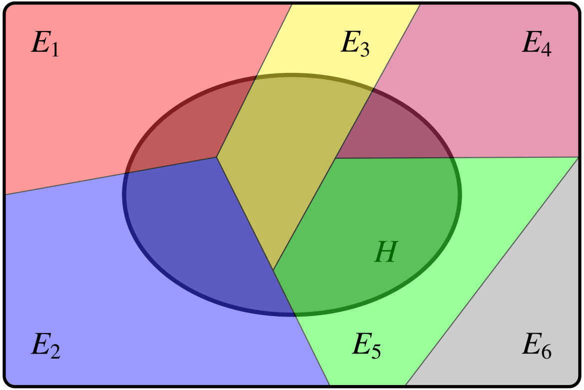

$$
\begin{aligned}
P(H) &= P(H \cap E_1) + \cdots + P(H \cap E_6) \\
&= \tfrac{P(E_1)}{P(E_1)} P(H \cap E_1) + \cdots + \tfrac{P(E_6)}{P(E_6)} P(H \cap E_6) \\
&= \textstyle\sum_i P(E_i) P(H \mid E_i)
\end{aligned}
$$

A basic result about probabilities.

For any finite partition $\mathcal{E} = \{E_i\}$ of F and and event H,

$$P(H \mid F) = \sum_i P(E_i \mid F) P(H \mid E_i)$$

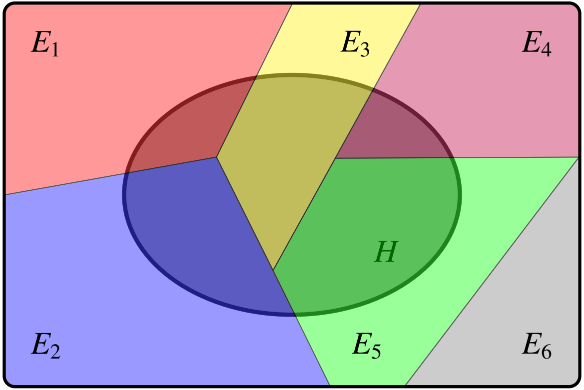

$$
\begin{aligned}
P(H \mid W) &= \textstyle\sum_i P(E_i \mid W) P(H \mid E_i \cap W) \\
&= \textstyle\sum_i P(E_i \mid W) P(H \mid E_i)
\end{aligned}
$$

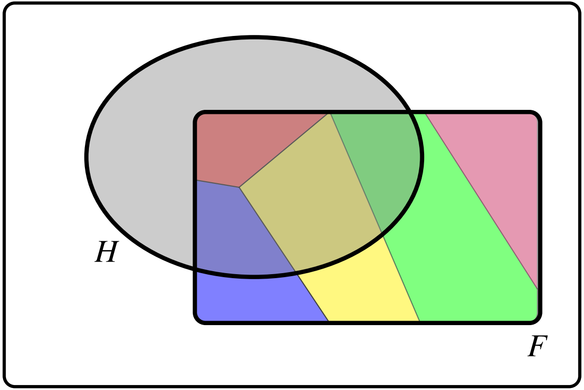

$$
\begin{aligned}
P(H \mid F) &= \textstyle\sum_i P(E_i \mid F) P(H \mid E_i \cap F) \\
&= \textstyle\sum_i P(E_i \mid F) P(H \mid E_i)
\end{aligned}
$$

Everyone Knows: $K(H) = \bigcap_{i \in \mathcal{A}} K_i(H)$

$K^m(H)$ for all $m \geq 0$ is defined as:

$K^0(H) = H \qquad K^m(H) = K(K^{m-1}(H))$

Common Knowledge: $C : \wp(W) \to \wp(W)$ with

$$C(H) = \bigcap_{m \geq 0} K^m(H)$$

$I_C(w) = \{v \mid \text{there is a finite path from } w \text{ to } v\}$

$C(H) = \{w \mid I_C(w) \subseteq H\}$

$I_C(w) = \{v \mid \text{there is a finite path from } w \text{ to } v\}$

$C(H) = \{w \mid I_C(w) \subseteq H\}$

Alternatively,
$w \in C(H)$ provided that there is a $F \subseteq W$ such that

1. $F \subseteq K(F)$
2. $F \subseteq H$

Theorem. Suppose that n agents share a common prior and have different private information. If there is common knowledge in the group of the posterior probabilities, then the posteriors must be equal.

Robert Aumann. *Agreeing to Disagree*. Annals of Statistics **4** (1976).

Suppose that W is, $E \subseteq W$ is an event, and two (or more) agents with partitions $\mathcal{E}_i$. Let P be the **common prior**.

The agent's posterior probabilities of the event E are *random variables*: $P_i^E : W \to [0, 1]$, $P_i^E(w) = P(E \mid \mathcal{E}_i[w])$.

So, $[\![P_i^E = r]\!] = \{w \mid P_i^E(w) = r\}$

Assume that $w \in C([\![P_1^E = r \wedge P_2^E = q]\!])$.

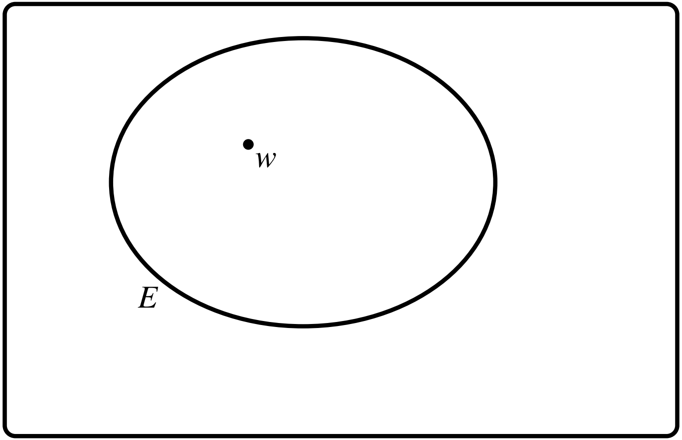

w
E

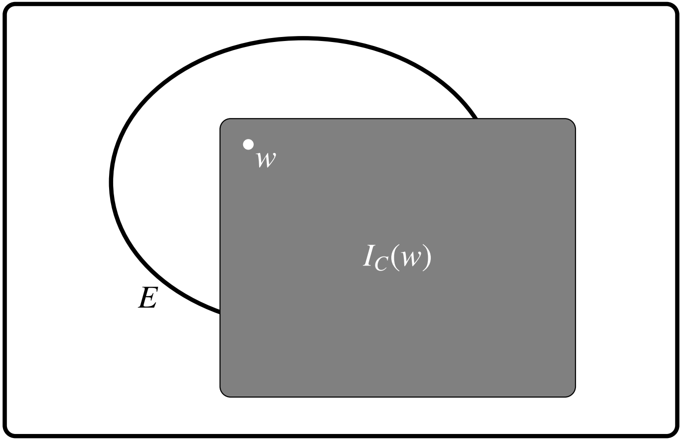

$$I_C(w) \subseteq [\![P_1^E = r \wedge P_2^E = q]\!]$$

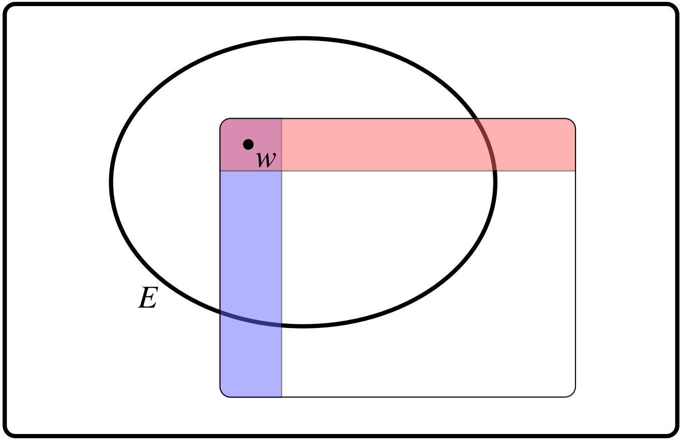

$$P(E \mid \mathcal{E}_1[w]) = q,\ P(E \mid \mathcal{E}_2[w]) = r$$

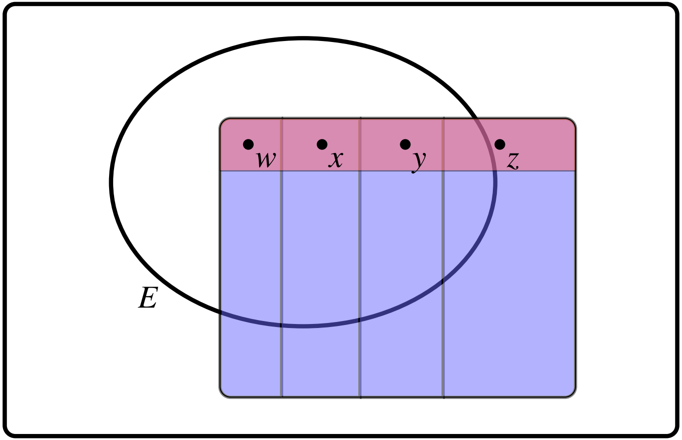

$$P(E \mid \mathcal{E}_1[w]) = P(E \mid \mathcal{E}_1[x]) = P(E \mid \mathcal{E}_1[y]) = P(E \mid \mathcal{E}_1[z]) = q$$

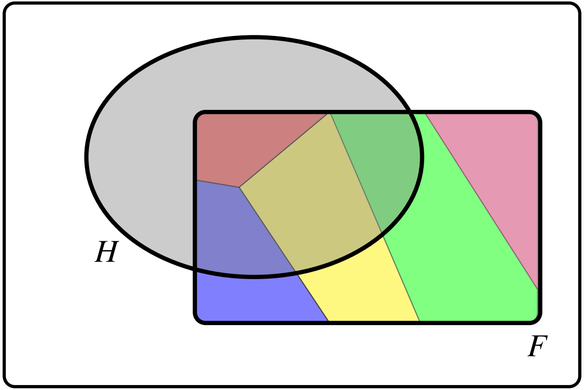

$$P(H \mid F) \quad = \quad \textstyle\sum_i P(E_i \mid F) P(H \mid E_i)$$

Fact. If $P(H \mid E_i) = q$ for all i, then $P(H \mid F) = q$.

Fact. Suppose that $\{F_i\}$ is a partition of F (so $F = \bigcup_i F_i$ and $F_i \cap F_j \neq \emptyset$ for $i \neq j$). If $P(E \mid F_i) = q$ for all i, then $P(E \mid F) = q$.

If $P(E \mid F_i) = q$, then $P(E \cap F_i) = qP(F_i)$.

$$P(E \mid F) = \frac{P(E \cap F)}{P(F)} = \frac{P((E \cap F_1) \cup \cdots \cup (E \cap F_n))}{P(F)}$$

$$= \frac{P(E \cap F_1) + \cdots + P(E \cap F_n)}{P(F)} = \frac{qP(F_1) + \cdots + qP(F_n)}{P(F)}$$

$$= \frac{q(P(F_1) + \cdots + P(F_n))}{P(F)} = \frac{qP(F)}{P(F)} = q$$

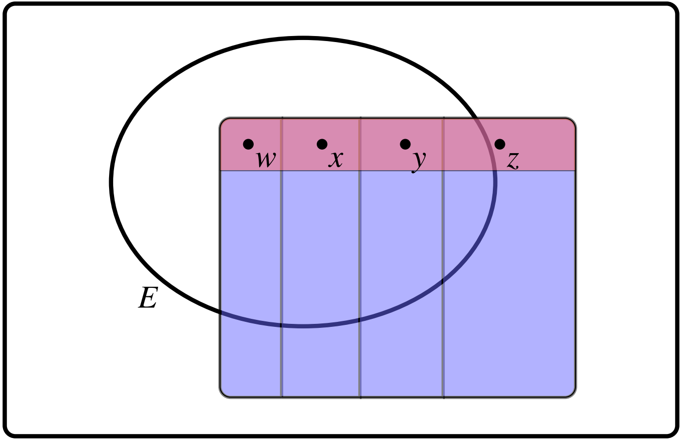

$$P(E \mid \mathcal{E}_1[w]) = P(E \mid \mathcal{E}_1[x]) = P(E \mid \mathcal{E}_1[y]) = P(E \mid \mathcal{E}_1[z]) = q$$

So, $P(E \mid I_C(w)) = q$.

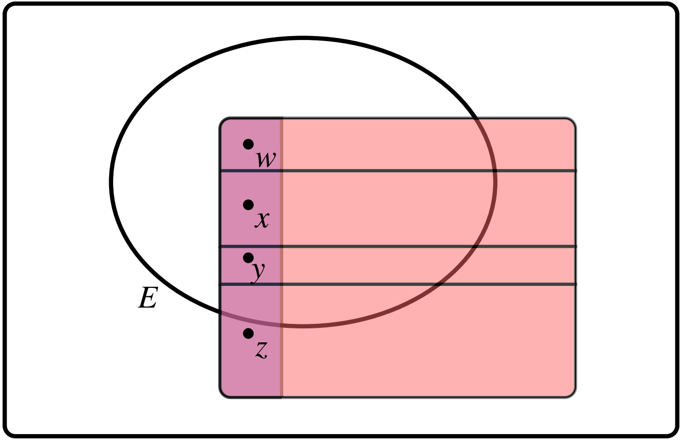

$$P(E \mid \mathcal{E}_2[w]) = P(E \mid \mathcal{E}_2[x]) = P(E \mid \mathcal{E}_2[y]) = P(E \mid \mathcal{E}_2[z]) = r$$

So, $P(E \mid I_C(w)) = r$.

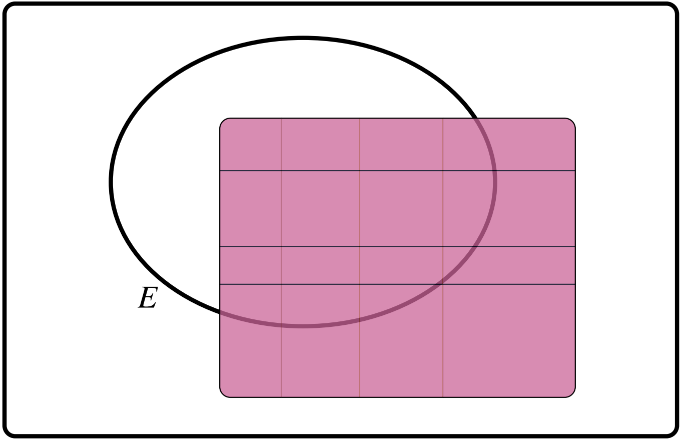

Thus, $q = P(E \mid I_C(w)) = r$.

Common r-belief

The typical example of an event that creates common knowledge is a **public announcement**.

Common r-belief

The typical example of an event that creates common knowledge is a **public announcement**.

Shouldn't one always allow for some small probability that a participant was absentminded, not listening, sending a text, checking Facebook, proving a theorem, asleep, ...

D. Monderer and D. Samet. *Approximating Common Knowledge with Common Beliefs*. Games and Economic Behavior (1989).

From Knowledge to r-Belief

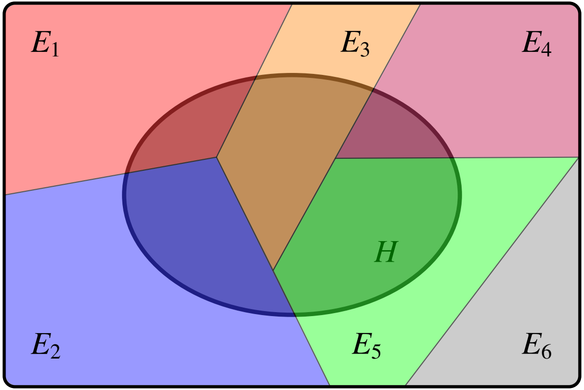

Given a partition $\mathcal{E}$, define $K_{\mathcal{E}} : \wp(W) \to \wp(W)$ as:

$$K_{\mathcal{E}}(H) = \{w \mid \mathcal{E}[w] \subseteq H\}$$

From Knowledge to r-Belief

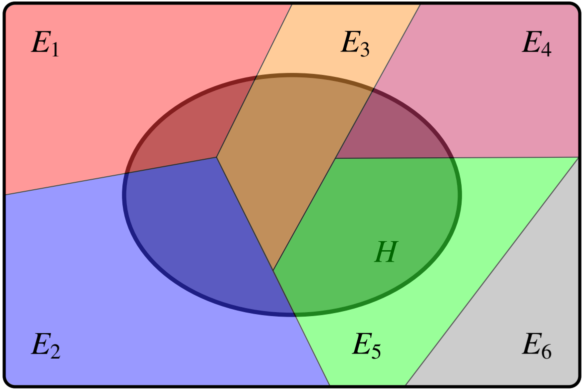

Given $r \in [0, 1]$ and a partition $\mathcal{E}$, define $B^r_{\mathcal{E}} : \wp(W) \rightarrow \wp(W)$
as: $B^r_{\mathcal{E}}(H) = \{w \mid P_{\mathcal{E},w}(H) \geq r\}$

From Knowledge to r-Belief

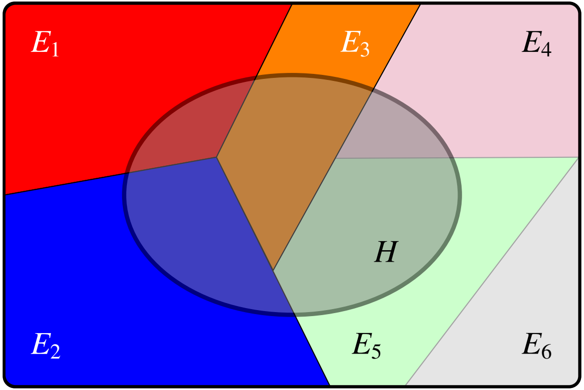

Given $r \in [0, 1]$ and a partition $\mathcal{E}$, define $B^r_{\mathcal{E}} : \wp(W) \to \wp(W)$
as: $B^r_{\mathcal{E}}(H) = \{w \mid P_{\mathcal{E},w}(H) \geq r\}$

From Common Knowledge to Common r-Belief

Suppose that $C : \wp(W) \to \wp(W)$ is a common knowledge operator. TFAE

1. $w \in C(H) = \bigcap_{m \geq 0} K^m(H)$
2. $I_c(w) \subseteq H$
3. There is a set $F \subseteq W$ such that
    3.1 $w \in F \subseteq K(F) = \bigcap_i K_i(F)$
    3.2 $F \subseteq H$

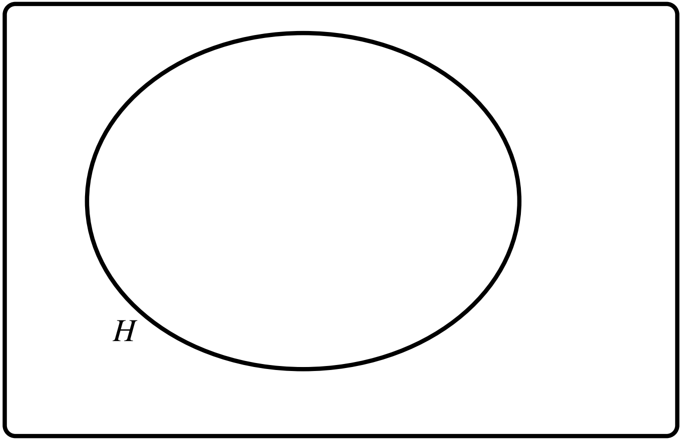
H

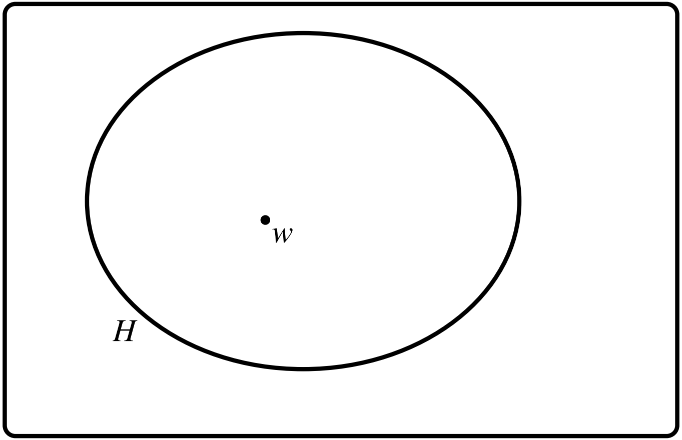
w
H

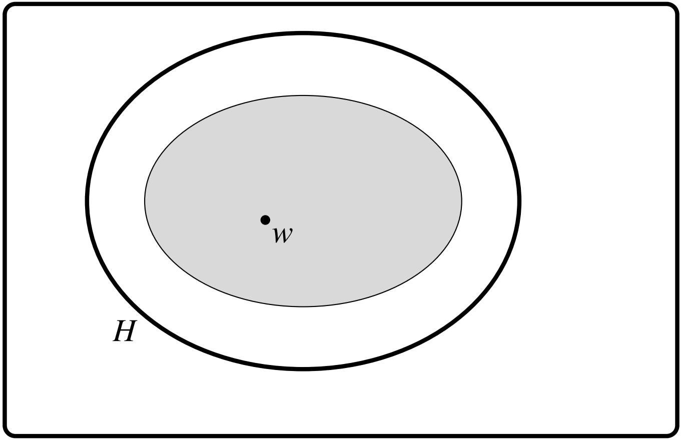
w
H

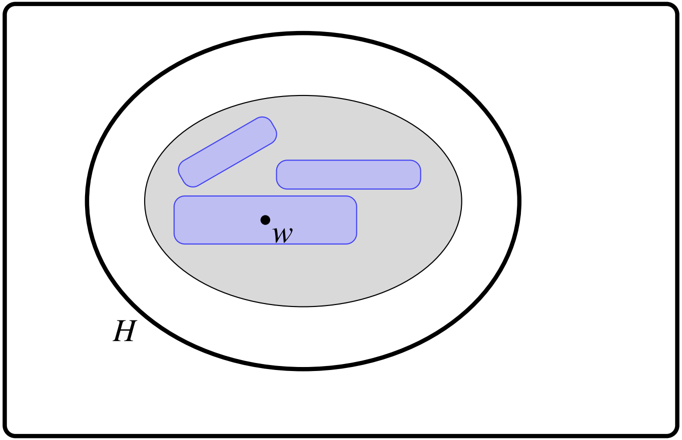
w
H

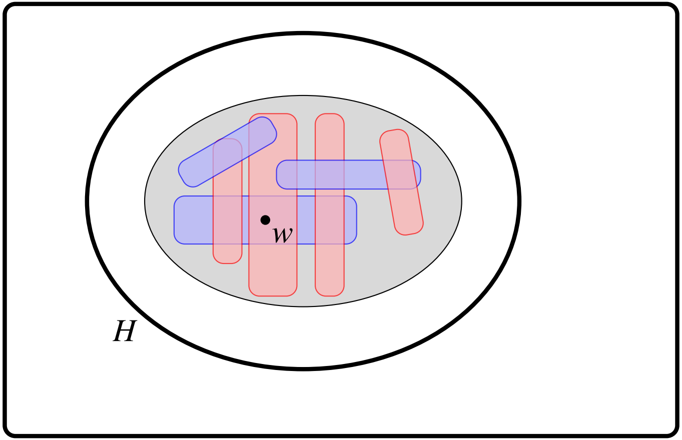
w
H

From Common Knowledge to Common r-Belief

$B_i^r(E) = \{w \mid P(E \mid \mathcal{E}_i[w]) \geq r\}$

From Common Knowledge to Common r-Belief

$B_i^r(E) = \{w \mid P(E \mid \mathcal{E}_i[w]) \geq r\}$

F is an **evident r-belief** if for each $i \in \mathcal{A}$, $F \subseteq B_i^r(F)$

From Common Knowledge to Common r-Belief

$B_i^r(E) = \{w \mid P(E \mid \mathcal{E}_i[w]) \geq r\}$

F is an **evident r-belief** if for each $i \in \mathcal{A}$, $F \subseteq B_i^r(F)$

An event H is **common r-belief** at w if there exists and evident r-belief event F such that $w \in F$ and for all $i \in \mathcal{A}$, $F \subseteq B_i^r(H)$

$w \in C(H)$ iff there is an event $F \subseteq W$ such that

1. $w \in F \subseteq K(F) = \bigcap_i K_i(F)$
2. $F \subseteq H$

$w \in C^r(H)$ iff there is an event $F \subseteq W$ such that

1. $w \in F \subseteq B^r(F) = \bigcap_i B_i^r(F)$
2. $F \subseteq B^r(H)$

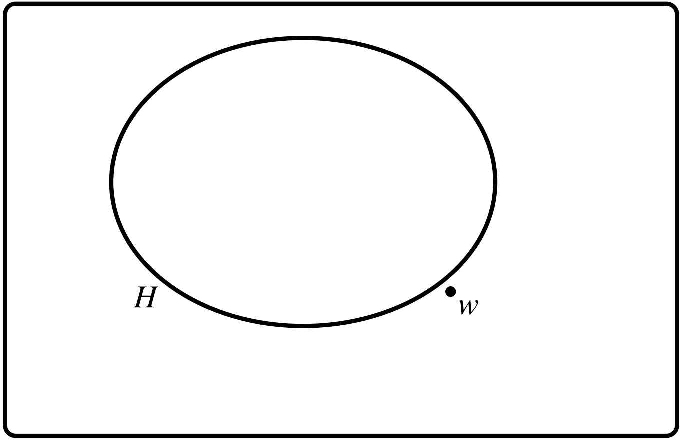
H
w

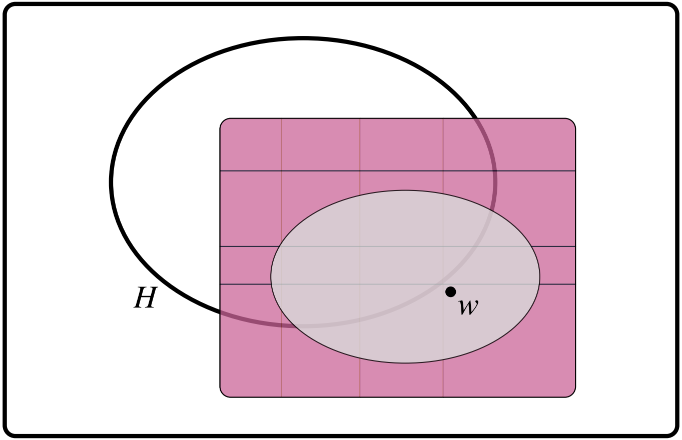
H
w

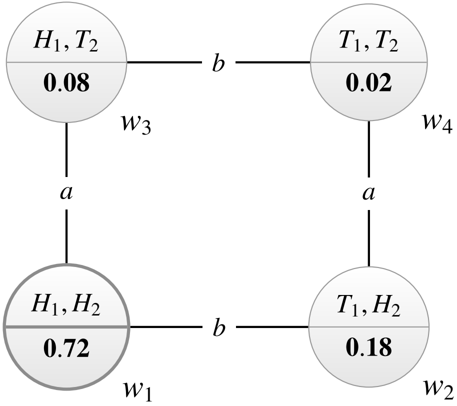

H_1, T_2
0.08
w_3
b
T_1, T_2
0.02
w_4
a
a
H_1, H_2
0.72
w_1
b
T_1, H_2
0.18
w_2

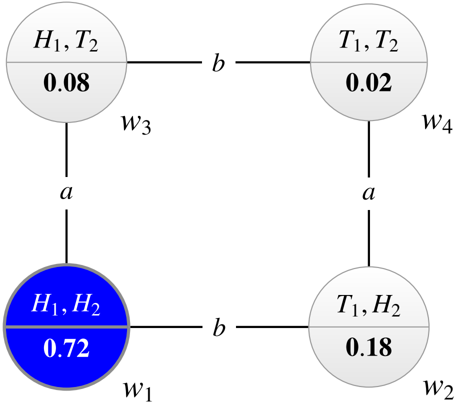

- $\{w_1\} = B_a^{0.9}(H_1 \cap H_2) \cap B_b^{0.8}(H_1 \cap H_2)$.
- $X = \{w_1\}$ is an evident 0.8-belief for both Ann and Bob.

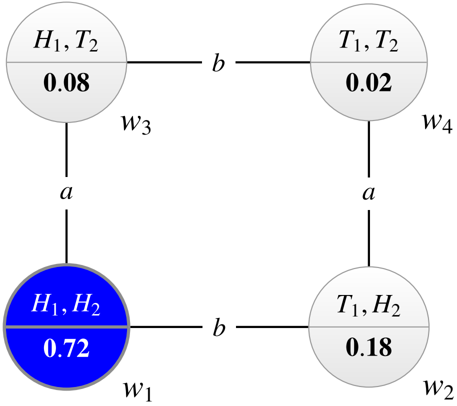

- $\{w_1\} = B_a^{0.9}(H_1 \cap H_2) \cap B_b^{0.8}(H_1 \cap H_2)$.
- $X = \{w_1\}$ is an evident 0.8-belief for both Ann and Bob.
- $X \subseteq B_a^{0.8}(H_1 \cap H_2) \cap B_b^{0.8}(H_1 \cap H_2)$.
- $w_1 \in C_{a,b}^{0.8}(H_1 \cap H_2)$.

Generalizing Aumann's Theorem

Theorem. If the posteriors of an event E are common p-belief at some state w, then any two posteriors can differ by at most $1 - p$.

D. Samet and D. Monderer. *Approximating Common Knowledge with Common Beliefs*. Games and Economic Behavior, Vol. 1, No. 2, 1989.

Assume that $w \in C^p([\![P_1^E = r \wedge P_2^E = q]\!])$.
There is an $F \subseteq W$ such that:

1. $F \subseteq B^p(F) = \bigcap_i B_i^p(F)$
2. $F \subseteq B^p([\![P_1^E = r \wedge P_2^E = q]\!]) = \bigcap_i B_i^p([\![P_1^E = r \wedge P_2^E = q]\!])$

Fact. For any H, Z_1, Z_2, $P(H \mid Z_1) \geq P(Z_2 \mid Z_1) P(H \mid Z_1 \cap Z_2)$

Fact. For any H, Z_1, Z_2, $P(H \mid Z_1) \geq P(Z_2 \mid Z_1) P(H \mid Z_1 \cap Z_2)$

$$P(H \mid Z_1) = \frac{P(H \cap Z_1)}{P(Z_1)}$$

Fact. For any H, Z_1, Z_2, $P(H \mid Z_1) \geq P(Z_2 \mid Z_1) P(H \mid Z_1 \cap Z_2)$

$$
\begin{aligned}
P(H \mid Z_1) &= \frac{P(H \cap Z_1)}{P(Z_1)} \\
&= \frac{P(Z_1 \cap Z_2)}{P(Z_1 \cap Z_2)} \frac{P(H \cap Z_1)}{P(Z_1)}
\end{aligned}
$$

Fact. For any H, Z_1, Z_2, $P(H \mid Z_1) \geq P(Z_2 \mid Z_1) P(H \mid Z_1 \cap Z_2)$

$$\begin{aligned} P(H \mid Z_1) &= \frac{P(H \cap Z_1)}{P(Z_1)} \\ &= \frac{P(Z_1 \cap Z_2)}{P(Z_1 \cap Z_2)} \frac{P(H \cap Z_1)}{P(Z_1)} \\ &= \frac{P(Z_1 \cap Z_2)}{P(Z_1)} \frac{P(H \cap Z_1)}{P(Z_1 \cap Z_2)} \end{aligned}$$

Fact. For any H, Z_1, Z_2, $P(H \mid Z_1) \geq P(Z_2 \mid Z_1) P(H \mid Z_1 \cap Z_2)$

$$\begin{aligned}
P(H \mid Z_1) &= \frac{P(H \cap Z_1)}{P(Z_1)} \\
&= \frac{P(Z_1 \cap Z_2)}{P(Z_1 \cap Z_2)} \frac{P(H \cap Z_1)}{P(Z_1)} \\
&= \frac{P(Z_1 \cap Z_2)}{P(Z_1)} \frac{P(H \cap Z_1)}{P(Z_1 \cap Z_2)} \\
&= P(Z_2 \mid Z_1) \frac{P(H \cap Z_1)}{P(Z_1 \cap Z_2)}
\end{aligned}$$

Fact. For any H, Z_1, Z_2, $P(H \mid Z_1) \geq P(Z_2 \mid Z_1) P(H \mid Z_1 \cap Z_2)$

$$
\begin{aligned}
P(H \mid Z_1) &= \frac{P(H \cap Z_1)}{P(Z_1)} \\
&= \frac{P(Z_1 \cap Z_2)}{P(Z_1 \cap Z_2)} \frac{P(H \cap Z_1)}{P(Z_1)} \\
&= \frac{P(Z_1 \cap Z_2)}{P(Z_1)} \frac{P(H \cap Z_1)}{P(Z_1 \cap Z_2)} \\
&= P(Z_2 \mid Z_1) \frac{P(H \cap Z_1)}{P(Z_1 \cap Z_2)} \\
&\geq P(Z_2 \mid Z_1) \frac{P(H \cap Z_1 \cap Z_2)}{P(Z_1 \cap Z_2)}
\end{aligned}
$$

Fact. For any H, Z_1, Z_2, $P(H \mid Z_1) \geq P(Z_2 \mid Z_1) P(H \mid Z_1 \cap Z_2)$

$$
\begin{aligned}
P(H \mid Z_1) &= \frac{P(H \cap Z_1)}{P(Z_1)} \\
&= \frac{P(Z_1 \cap Z_2)}{P(Z_1 \cap Z_2)} \frac{P(H \cap Z_1)}{P(Z_1)} \\
&= \frac{P(Z_1 \cap Z_2)}{P(Z_1)} \frac{P(H \cap Z_1)}{P(Z_1 \cap Z_2)} \\
&= P(Z_2 \mid Z_1) \frac{P(H \cap Z_1)}{P(Z_1 \cap Z_2)} \\
&\geq P(Z_2 \mid Z_1) \frac{P(H \cap Z_1 \cap Z_2)}{P(Z_1 \cap Z_2)} \\
&= P(Z_2 \mid Z_1) P(H \mid Z_1 \cap Z_2)
\end{aligned}
$$

Assume that $w \in C^p([\![P_1^E = r \wedge P_2^E = q]\!])$. There is an $F \subseteq W$ such that:

1. $F \subseteq B^p(F) = \bigcap_i B_i^p(F)$
2. $F \subseteq B^p([\![P_1^E = r \wedge P_2^E = q]\!]) = \bigcap_i B_i^p([\![P_1^E = r \wedge P_2^E = q]\!])$

Let $Z_1 = B_1^p(F)$ and $Z_2 = B_2^p(F)$.

Assume that $w \in C^p([\![P_1^E = r \wedge P_2^E = q]\!])$. There is an $F \subseteq W$ such that:

1. $F \subseteq B^p(F) = \bigcap_i B_i^p(F)$
2. $F \subseteq B^p([\![P_1^E = r \wedge P_2^E = q]\!]) = \bigcap_i B_i^p([\![P_1^E = r \wedge P_2^E = q]\!])$

Let $Z_1 = B_1^p(F)$ and $Z_2 = B_2^p(F)$.

From the previous Fact:

1. $P(E \mid Z_1) \geq P(Z_2 \mid Z_1)P(E \mid Z_1 \cap Z_2)$
2. $P(\overline{E} \mid Z_1) \geq P(Z_2 \mid Z_1)P(\overline{E} \mid Z_1 \cap Z_2)$

Assume that $w \in C^p(\llbracket P_1^E = r \wedge P_2^E = q \rrbracket)$. There is an $F \subseteq W$ such that:

1. $F \subseteq B^p(F) = \bigcap_i B_i^p(F)$
2. $F \subseteq B^p(\llbracket P_1^E = r \wedge P_2^E = q \rrbracket) = \bigcap_i B_i^p(\llbracket P_1^E = r \wedge P_2^E = q \rrbracket)$

Let $Z_1 = B_1^p(F)$ and $Z_2 = B_2^p(F)$.

Since $P(Z_2 \mid Z_1) \geq P(B^p(E) \mid Z_1) \geq p$,

1. $P(E \mid Z_1) \geq P(Z_2 \mid Z_1)P(E \mid Z_1 \cap Z_2) \geq pP(E \mid Z_1 \cap Z_2)$
2. $P(\overline{E} \mid Z_1) \geq P(Z_2 \mid Z_1)P(\overline{E} \mid Z_1 \cap Z_2) \geq pP(\overline{E} \mid Z_1 \cap Z_2)$

Assume that $w \in C^p([\![P_1^E = r \wedge P_2^E = q]\!])$. There is an $F \subseteq W$ such that:

1. $F \subseteq B^p(F) = \bigcap_i B_i^p(F)$
2. $F \subseteq B^p([\![P_1^E = r \wedge P_2^E = q]\!]) = \bigcap_i B_i^p([\![P_1^E = r \wedge P_2^E = q]\!])$

Let $Z_1 = B_1^p(F)$ and $Z_2 = B_2^p(F)$.

Since $P(Z_2 \mid Z_1) \geq P(B^p(E) \mid Z_1) \geq p$,

1. $P(E \mid Z_1) \geq P(Z_2 \mid Z_1)P(E \mid Z_1 \cap Z_2) \geq pP(E \mid Z_1 \cap Z_2)$
2. $P(\overline{E} \mid Z_1) \geq P(Z_2 \mid Z_1)P(\overline{E} \mid Z_1 \cap Z_2) \geq pP(\overline{E} \mid Z_1 \cap Z_2)$
   So, $1 - P(E \mid Z_1) \geq p(1 - P(E \mid Z_1 \cap Z_2))$

Assume that $w \in C^p([\![P_1^E = r \wedge P_2^E = q]\!])$. There is an $F \subseteq W$ such that:

1. $F \subseteq B^p(F) = \bigcap_i B_i^p(F)$
2. $F \subseteq B^p([\![P_1^E = r \wedge P_2^E = q]\!]) = \bigcap_i B_i^p([\![P_1^E = r \wedge P_2^E = q]\!])$

Let $Z_1 = B_1^p(F)$ and $Z_2 = B_2^p(F)$.

Since $P(E \mid Z_1) = r$,

1. $P(E \mid Z_1) \geq pP(E \mid Z_1 \cap Z_2)$
   So, $r \geq pP(E \mid Z_1 \cap Z_2)$
2. $1 - P(E \mid Z_1) \geq p(1 - P(E \mid Z_1 \cap Z_2))$
   So, $1 - r \geq p(1 - P(E \mid Z_1 \cap Z_2))$

Assume that $w \in C^p([\![P_1^E = r \wedge P_2^E = q]\!])$. There is an $F \subseteq W$ such that:

1. $F \subseteq B^p(F) = \bigcap_i B_i^p(F)$
2. $F \subseteq B^p([\![P_1^E = r \wedge P_2^E = q]\!]) = \bigcap_i B_i^p([\![P_1^E = r \wedge P_2^E = q]\!])$

Let $Z_1 = B_1^p(F)$ and $Z_2 = B_2^p(F)$.

Since $P(E \mid Z_1) = r$,

1. $P(E \mid Z_1) \geq pP(E \mid Z_1 \cap Z_2)$
   So, $r \geq pP(E \mid Z_1 \cap Z_2)$
2. $1 - P(E \mid Z_1) \geq p(1 - P(E \mid Z_1 \cap Z_2))$
   So, $1 - r \geq p(1 - P(E \mid Z_1 \cap Z_2))$

$$pP(E \mid Z_1 \cap Z_2) \leq r \leq 1 - p + pP(E \mid Z_1 \cap Z_2)$$

Assume that $w \in C^p([\![P_1^E = r \wedge P_2^E = q]\!])$. There is an $F \subseteq W$ such that:

1. $F \subseteq B^p(F) = \bigcap_i B_i^p(F)$
2. $F \subseteq B^p([\![P_1^E = r \wedge P_2^E = q]\!]) = \bigcap_i B_i^p([\![P_1^E = r \wedge P_2^E = q]\!])$

Let $Z_1 = B_1^p(F)$ and $Z_2 = B_2^p(F)$.

(Similar argument for player 2: $P(E \mid Z_2) = r$ and $P(Z_1 \mid Z_2) \geq p$)

$$pP(E \mid Z_1 \cap Z_2) \leq r \leq 1 - p + pP(E \mid Z_1 \cap Z_2)$$

$$pP(E \mid Z_2 \cap Z_1) \leq q \leq 1 - p + pP(E \mid Z_2 \cap Z_1)$$

Assume that $w \in C^p([\![P_1^E = r \wedge P_2^E = q]\!])$. There is an $F \subseteq W$ such that:

1. $F \subseteq B^p(F) = \bigcap_i B_i^p(F)$
2. $F \subseteq B^p([\![P_1^E = r \wedge P_2^E = q]\!]) = \bigcap_i B_i^p([\![P_1^E = r \wedge P_2^E = q]\!])$

Let $Z_1 = B_1^p(F)$ and $Z_2 = B_2^p(F)$.

(Similar argument for player 2: $P(E \mid Z_2) = r$ and $P(Z_1 \mid Z_2) \geq p$)

$$pP(E \mid Z_1 \cap Z_2) \leq r \leq 1 - p + pP(E \mid Z_1 \cap Z_2)$$

$$pP(E \mid Z_2 \cap Z_1) \leq q \leq 1 - p + pP(E \mid Z_2 \cap Z_1)$$

Hence, $|r - q| \leq 1 - p + pP(E \mid Z_2 \cap Z_1) - pP(E \mid Z_2 \cap Z_1) = 1 - p$

Dynamic characterization of Aumann's Theorem

- How do the posteriors *become* common knowledge?

J. Geanakoplos and H. Polemarchakis. *We Can't Disagree Forever*. Journal of Economic Theory (1982).

Dynamic characterization of Aumann's Theorem

- How do the posteriors *become* common knowledge?

J. Geanakoplos and H. Polemarchakis. *We Can't Disagree Forever*. Journal of Economic Theory (1982).

- What happens when communication is not the the whole group, but pairwise?

R. Parikh and P. Krasucki. *Communication, Consensus and Knowledge*. Journal of Economic Theory (1990).

$$t = 0 \qquad \langle W, \mathcal{E}_{0,a}, \mathcal{E}_{0,b}, p \rangle$$

$$t = 0 \qquad \langle W, \mathcal{E}_{0,a}, \mathcal{E}_{0,b}, p \rangle$$

$$P^E_{0,a}(w) = r_0 \quad P^E_{0,b}(w) = q_0$$

$t = 0 \qquad \langle W, \mathcal{E}_{0,a}, \mathcal{E}_{0,b}, p \rangle$

$$P^E_{0,a}(w) = r_0 \quad P^E_{0,b}(w) = q_0$$

$t = 1 \qquad \langle W, \mathcal{E}_{1,a}, \mathcal{E}_{1,b}, p \rangle$

$t = 0 \qquad \langle W, \mathcal{E}_{0,a}, \mathcal{E}_{0,b}, p\rangle$

$$P^E_{0,a}(w) = r_0 \quad P^E_{0,b}(w) = q_0$$

$t = 1 \qquad \langle W, \mathcal{E}_{1,a}, \mathcal{E}_{1,b}, p\rangle$

$$P^E_{1,a}(w) = r_1 \quad P^E_{1,b}(w) = q_1$$

$t = 0 \qquad \langle W, \mathcal{E}_{0,a}, \mathcal{E}_{0,b}, p \rangle$

$$P^E_{0,a}(w) = r_0 \quad P^E_{0,b}(w) = q_0$$

$t = 1 \qquad \langle W, \mathcal{E}_{1,a}, \mathcal{E}_{1,b}, p \rangle$

$$P^E_{1,a}(w) = r_1 \quad P^E_{1,b}(w) = q_1$$

$t = 2 \qquad \langle W, \mathcal{E}_{2,a}, \mathcal{E}_{2,b}, p \rangle$

$$P^E_{2,a}(w) = r_2 \quad P^E_{2,b}(w) = q_2$$

$t = 3 \qquad \langle W, \mathcal{E}_{3,a}, \mathcal{E}_{3,b}, p \rangle$

$\vdots \qquad \vdots$

Geanakoplos and Polemarchakis

- Assuming that the information partitions are finite, given an event A, the revision process converges in finitely many steps.

Geanakoplos and Polemarchakis

- Assuming that the information partitions are finite, given an event A, the revision process converges in finitely many steps.

- For each n, there are examples where the process takes n steps.

Geanakoplos and Polemarchakis

- Assuming that the information partitions are finite, given an event A, the revision process converges in finitely many steps.
- For each n, there are examples where the process takes n steps.
- An *indirect communication* equilibrium is not necessarily a *direct communication* equilibrium.

What type of information exchanges should be used in a dynamic characterization of Monderer and Samet's generalization of Aumann's Theorem?

That is, for an event F and an epistemic-probability model, what dynamic process will converge on a model in which there is common p-belief of the agents' current probabilities of F?