Sheaves of Probability

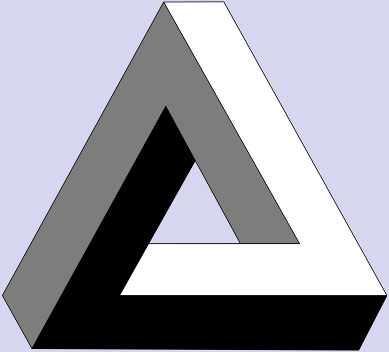

Owen Biesel

Southern Connecticut State University
bieselo1@southernct.edu

Credences represented by a probabilities

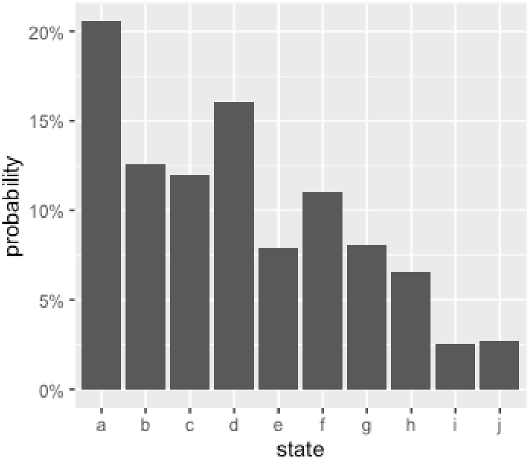

Our credences are represented by a probability measure P on a "state space," e.g. $X = \{a, b, c, d, e, f, g, h, i, j\}$.

Conditionalization

$P:$

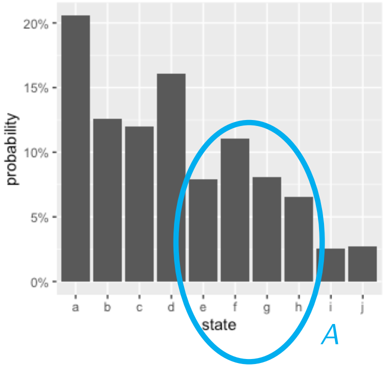

$\rightsquigarrow P|_A :$

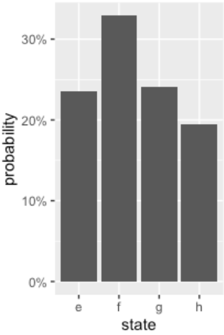

Gradually accumulating evidence narrows down state space

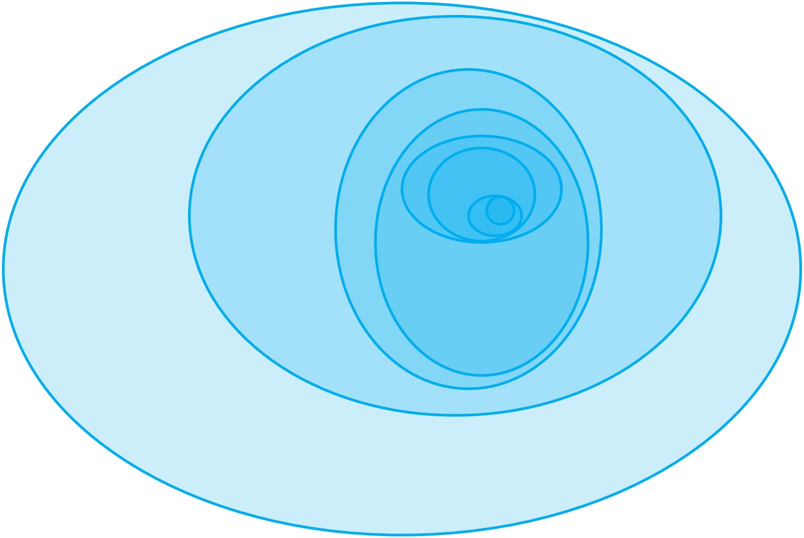

Do we actually accumulate evidence monotonically?

Show of hands:

- How many of you can no longer remember something you're sure you used to know?

Overlapping sets of evidence

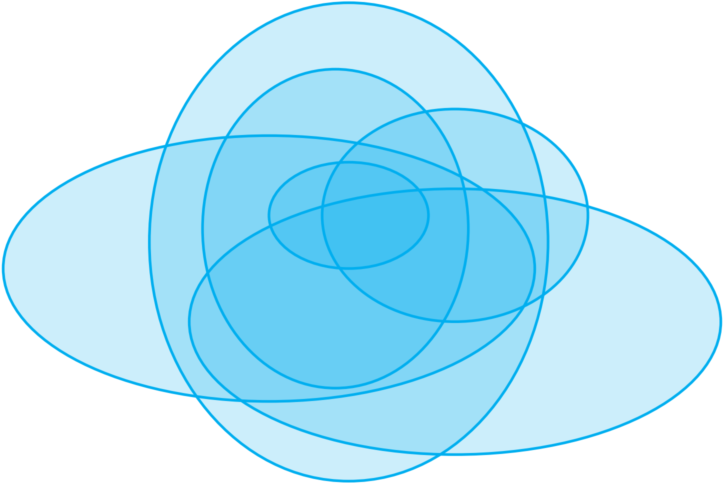

Generalized Conditionalization

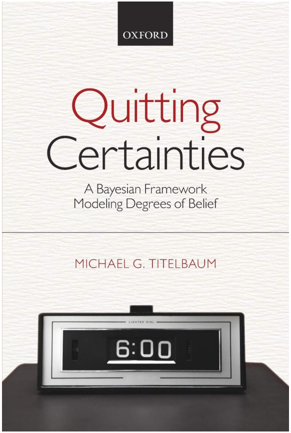

Titelbaum's notion of consistency among a collection of different agents' credences:

Definition (Generalized Conditionalization)

Let X be a state space, and let P_i be probability measures on subsets $A_i \subseteq X$. We say that the P_i satisfy *generalized conditionalization* (GC) if for each pair of agents i, j we have

$$P_i|_{A_i \cap A_j} = P_j|_{A_i \cap A_j}.$$

Essentially, do the P_i agree where they overlap?

Example credences satisfying GC

P_1:

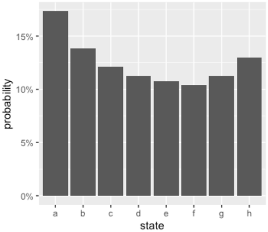

P_2:

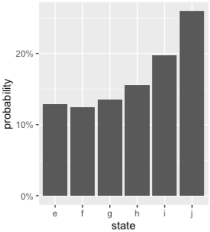

$$P_1|_{A_1 \cap A_2} =$$

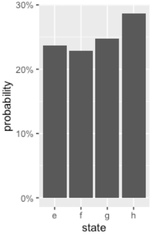

$$= P_2|_{A_1 \cap A_2}$$

A common prior

P:

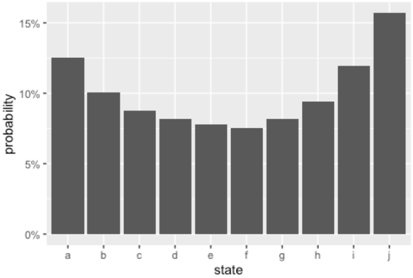

$P_1 = P|_{A_1}$:

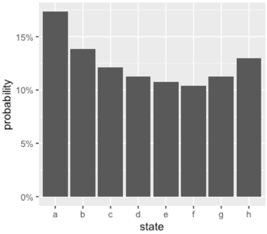

$P_2 = P|_{A_2}$:

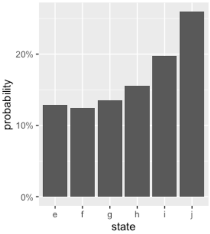

Sheaves

If probability measures P_i on A_i satisfying GC (namely, $P_i|_{A_i \cap A_j} = P_j|_{A_i \cap A_j}$) always uniquely determine a "common prior" P on $\bigcup_i A_i$ with $P|_{A_i} = P_i$, it would mean that probability measures form a *sheaf*. Do they?

Sheaves

If probability measures P_i on A_i satisfying GC (namely, $P_i|_{A_i \cap A_j} = P_j|_{A_i \cap A_j}$) always uniquely determine a "common prior" P on $\bigcup_i A_i$ with $P|_{A_i} = P_i$, it would mean that probability measures form a *sheaf*. Do they?

Answer: No!

GC credences with no common prior

	a	b	c
P_1	40%	60%	
P_2		40%	60%
P_3	60%		40%

These have no common prior P: we would need to have

$$P(a) < P(b) < P(c) < P(a).$$

They're also "logically inconsistent" in the sense that together, the three agents have learned evidence that rules out all three possible states.

Verifying GC

	a	b
P_1	40%	60%

	a	c
P_3	60%	40%

	b	c
P_2	40%	60%

Verifying GC

	a	b
P_1	40%	60%

a
100%

	a	c
P_3	60%	40%

	b	c
P_2	40%	60%

Verifying GC

$$
\begin{tikzcd}
 & & \begin{array}{c|c|c} & a & b \\ \hline P_1 & 40\% & 60\% \end{array} \arrow[dl, squiggly] \arrow[dr, squiggly] & & \\
 & \begin{array}{c} a \\ \hline 100\% \end{array} & & \begin{array}{c} b \\ \hline 100\% \end{array} & \\
\begin{array}{c|c|c} & a & c \\ \hline P_3 & 60\% & 40\% \end{array} \arrow[ur, squiggly] \arrow[rr, squiggly] & & \begin{array}{c} c \\ \hline 100\% \end{array} & & \begin{array}{c|c|c} & b & c \\ \hline P_2 & 40\% & 60\% \end{array} \arrow[ul, squiggly] \arrow[ll, squiggly]
\end{tikzcd}
$$

Verifying GC

	a	b
P_1	40%	60%

a
100%

b
100%

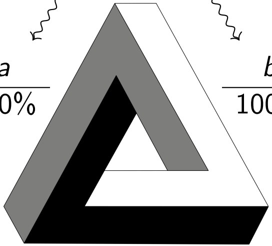

	a	c
P_3	60%	40%

c
100%

	b	c
P_2	40%	60%

Three notions of compatibility

1. GC (Generalized Conditionalization): Do the probability measures agree on their overlaps?
2. CP (Common Prior): Is each probability measure the restriction of a single prior distribution?
3. LC (Logical Consistency): Do the probability measures all overlap nontrivially?

We have seen an example where GC holds but CP and LC do not.

LC $\implies$ (GC $\iff$ CP)

In fact, for logically consistent probabilities, satisfying GC is equivalent to having a common prior!

Theorem (B—, '24)

Let X be a state space, and for each $i \in \{1, \ldots, n\}$ let P_i be a probability measure on $A_i \subseteq X$. Suppose that for each i, we have $P_i(\bigcap_{i=1}^n A_i) > 0$. Then the following are equivalent:

- *For all $i, j \in \{1, \ldots, n\}$, we have $P_i|_{A_i \cap A_j} = P_j|_{A_i \cap A_j}$.*
- *There exists a unique probability measure P on $\bigcup_{i=1}^n A_i$ such that for each $i \in \{1, \ldots, n\}$, we have $P|_{A_i} = P_i$.*

Sheaves of Probability

Theorem (B—, '24)

Let X be a measurable space, let $E \subseteq X$ be a measurable subset, and let $[E, X]$ be the collection of measurable subsets of X containing E, partially ordered by inclusion. Equip $[E, X]$ with the notion of coverage given by finite unions, making it a site. Then the functor

$$[E, X]^{\mathrm{op}} \to \mathbf{Set}$$

sending each subset A to $\{$probability measures P on A with $P(E) > 0\}$, and each inclusion $A \subseteq B$ to the restriction function $P \mapsto P|_A$, is a sheaf.

Sketch of Proof

Proof Sketch.

1. Scale each P_i to an ordinary finite measure μ_i such that $\mu_i(E) = 1$.
2. Show that the measures μ_i agree "on the nose" on their overlaps.
3. Construct a finite measure μ on $\bigcup_{i=1}^n A_i$ that agrees with each μ_i.
4. Normalize μ to give the desired probability measure P on $\bigcup_{i=1}^n A_i$.

$\square$

Thank you!

Questions for further discussion:

- What goes wrong with infinite covers $A_1, A_2, A_3, \ldots$?
- Is there some kind of group cohomology lurking in this story?
- Could we consider sheaves on more general categories to account for self-locating uncertainty?