Logical and Probabilistic Models of Belief Change

Eric Pacuit

Department of Philosophy
University of Maryland, College Park
pacuit.org

July 12, 2016

Plan

- Day 1 Introduction to belief revision, AGM, possible worlds models, Bayesian models (time permitted)
- Day 2 Bayesian models (continued), Justifying Bayesian models (Dutch books, Accuracy-based arguments), Updating probabilities
- Day 3 The value of learning, Lottery Paradox, Preface Paradox, Review Paradox, Iterated belief revision, Context shifts, Becoming aware
- Day 4 The value of learning, Lottery Paradox, Preface Paradox, Review Paradox, Iterated belief revision, Context shifts, Becoming aware (continued)
- Day 5 Interactive epistemology (Agreement Theorems, Belief Revision in Games)

pacuit.org/nasslli2016/belrev/

Plan for today

- Quick recap (AGM, possible worlds models)
- Bayesian models
-
- Updating probabilities

Dynamic Epistemic Logic

The key idea of dynamic epistemic logic is that we can represent changes in agents' epistemic states by *transforming models*.

Dynamic Epistemic Logic

The key idea of dynamic epistemic logic is that we can represent changes in agents' epistemic states by *transforming models*.

In the simplest case, we model an agent's acquisition of knowledge by the elimination of possibilities from an initial epistemic model.

Finding out that φ

$$\mathcal{M} = \langle W, \{\sim_i\}_{i\in\mathcal{A}}, \{\preceq_i\}_{i\in\mathcal{A}}, V\rangle$$

$$\Big\Downarrow \; \textbf{Find out that } \varphi \; \Big\Downarrow$$

$$\mathcal{M}' = \langle W', \{\sim_i'\}_{i\in\mathcal{A}}, \{\preceq_i'\}_{i\in\mathcal{A}}, V|_{W'}\rangle$$

Example: College Park and Amsterdam

Recall the College Park agent who doesn't know whether it's raining in Amsterdam, whose epistemic state is represented by the model:

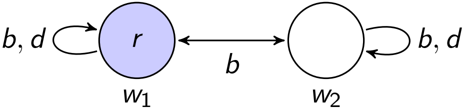

Example: College Park and Amsterdam

Recall the College Park agent who doesn't know whether it's raining in Amsterdam, whose epistemic state is represented by the model:

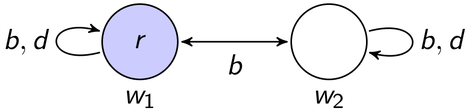

What happens when the Amsterdam agent calls the College Park agent on the phone and says, "It's raining in Amsterdam"?

Example: College Park and Amsterdam

Recall the College Park agent who doesn't know whether it's raining in Amsterdam, whose epistemic state is represented by the model:

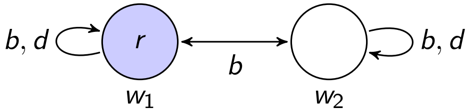

What happens when the Amsterdam agent calls the College Park agent on the phone and says, "It's raining in Amsterdam"?

We model the change in b's epistemic state by eliminating all epistemic possibilities in which it's *not* raining in Amsterdam.

Example: College Park and Amsterdam

Recall the College Park agent who doesn't know whether it's raining in Amsterdam, whose epistemic state is represented by the model:

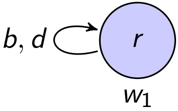

What happens when the Amsterdam agent calls the College Park agent on the phone and says, "It's raining in Amsterdam"?

We model the change in b's epistemic state by eliminating all epistemic possibilities in which it's *not* raining in Amsterdam.

Model Update

We can easily give a formal definition that captures the idea of knowledge acquisition as the elimination of possibilities.

Model Update

We can easily give a formal definition that captures the idea of knowledge acquisition as the elimination of possibilities.

Given $\mathcal{M} = \langle W, \{R_a \mid a \in \text{Agt}\}, V\rangle$, the *updated model* $\mathcal{M}_{|\varphi}$ is obtained by deleting from $\mathcal{M}$ all worlds in which φ was false.

Model Update

We can easily give a formal definition that captures the idea of knowledge acquisition as the elimination of possibilities.

Given $\mathcal{M} = \langle W, \{R_a \mid a \in \mathsf{Agt}\}, V\rangle$, the *updated model* $\mathcal{M}_{|\varphi}$ is obtained by deleting from $\mathcal{M}$ all worlds in which φ was false.

Formally, $\mathcal{M}_{|\varphi} = \langle W_{|\varphi}, \{R_{a_{|\varphi}} \mid a \in \mathsf{Agt}\}, V_{|\varphi}\rangle$ is the model s.th.:

$$W_{|\varphi} = \{v \in W \mid \mathcal{M}, v \models \varphi\};$$

Model Update

We can easily give a formal definition that captures the idea of knowledge acquisition as the elimination of possibilities.

Given $\mathcal{M} = \langle W, \{R_a \mid a \in \text{Agt}\}, V\rangle$, the *updated model* $\mathcal{M}_{|\varphi}$ is obtained by deleting from $\mathcal{M}$ all worlds in which φ was false.

Formally, $\mathcal{M}_{|\varphi} = \langle W_{|\varphi}, \{R_{a_{|\varphi}} \mid a \in \text{Agt}\}, V_{|\varphi}\rangle$ is the model s.th.:

$W_{|\varphi} = \{v \in W \mid \mathcal{M}, v \models \varphi\}$;

$R_{a_{|\varphi}}$ is the restriction of R_a to $W_{|\varphi}$;

Model Update

We can easily give a formal definition that captures the idea of knowledge acquisition as the elimination of possibilities.

Given $\mathcal{M} = \langle W, \{R_a \mid a \in \text{Agt}\}, V\rangle$, the *updated model* $\mathcal{M}_{|\varphi}$ is obtained by deleting from $\mathcal{M}$ all worlds in which φ was false.

Formally, $\mathcal{M}_{|\varphi} = \langle W_{|\varphi}, \{R_{a_{|\varphi}} \mid a \in \text{Agt}\}, V_{|\varphi}\rangle$ is the model s.th.:

$W_{|\varphi} = \{v \in W \mid \mathcal{M}, v \models \varphi\}$;

$R_{a_{|\varphi}}$ is the restriction of R_a to $W_{|\varphi}$;

$V_{|\varphi}(p)$ is the intersection of $V(p)$ and $W_{|\varphi}$.

Model Update

We can easily give a formal definition that captures the idea of knowledge acquisition as the elimination of possibilities.

Given $\mathcal{M} = \langle W, \{R_a \mid a \in \text{Agt}\}, V\rangle$, the *updated model* $\mathcal{M}_{|\varphi}$ is obtained by deleting from $\mathcal{M}$ all worlds in which φ was false.

Formally, $\mathcal{M}_{|\varphi} = \langle W_{|\varphi}, \{R_{a_{|\varphi}} \mid a \in \text{Agt}\}, V_{|\varphi}\rangle$ is the model s.th.:

$W_{|\varphi} = \{v \in W \mid \mathcal{M}, v \models \varphi\}$;

$R_{a_{|\varphi}}$ is the restriction of R_a to $W_{|\varphi}$;

$V_{|\varphi}(p)$ is the intersection of $V(p)$ and $W_{|\varphi}$.

In the single-agent case, this models the agent learning φ. In the multi-agent case, this models all agents *publicly* learning φ.

Public Announcement Logic

One of the **big ideas** of dynamic epistemic logic is to add to our formal language operators that can describe the kinds of model updates that we just saw for the College Park and Amsterdam example.

Public Announcement Logic

One of the **big ideas** of dynamic epistemic logic is to add to our formal language operators that can describe the kinds of model updates that we just saw for the College Park and Amsterdam example.

The language of Public Announcement Logic (PAL) is given by:

$$\varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid K_a\varphi \mid [!\varphi]\varphi$$

Public Announcement Logic

One of the **big ideas** of dynamic epistemic logic is to add to our formal language operators that can describe the kinds of model updates that we just saw for the College Park and Amsterdam example.

The language of Public Announcement Logic (PAL) is given by:

$$\varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid K_a\varphi \mid [!\varphi]\varphi$$

Read $[!\varphi]\psi$ as "after (every) true announcement of φ, ψ."

Public Announcement Logic

One of the **big ideas** of dynamic epistemic logic is to add to our formal language operators that can describe the kinds of model updates that we just saw for the College Park and Amsterdam example.

The language of Public Announcement Logic (PAL) is given by:

$$\varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid K_a\varphi \mid [!\varphi]\varphi$$

Read $[!\varphi]\psi$ as "after (every) true announcement of φ, ψ."

Read $\langle !\varphi\rangle\psi := \neg[!\varphi]\neg\psi$ as "after a true announcement of φ, ψ."

Public Announcement Logic

The truth clause for the dynamic operator $[!\varphi]$ is:

- $\mathcal{M}, w \models [!\varphi]\psi$ iff $\mathcal{M}, w \models \varphi$ *implies* $\mathcal{M}_{|\varphi}, w \models \psi$.

  So if φ is false, $[!\varphi]\psi$ is vacuously true. Here is the $\langle !\varphi \rangle$ clause:

- $\mathcal{M}, w \models \langle !\varphi \rangle\psi$ iff $\mathcal{M}, w \models \varphi$ *and* $\mathcal{M}_{|\varphi}, w \models \psi$.

Main Idea: we evaluate $[!\varphi]\psi$ and $\langle !\varphi \rangle\psi$ not by looking at *other worlds in the same model*, but rather by looking at a new model.

Public Announcement Logic

Suppose $\mathcal{M} = \langle W, \{\sim_i\}_{i\in\mathcal{A}}, \{\preceq_i\}_{i\in\mathcal{A}}, V\rangle$ is a multi-agent Kripke Model

$$\mathcal{M}, w \models [!\psi]\varphi \text{ iff } \mathcal{M}, w \models \psi \text{ implies } \mathcal{M}|_\psi, w \models \varphi$$

where $\mathcal{M}|_\psi = \langle W', \{\sim'_i\}_{i\in\mathcal{A}}, \{\preceq'_i\}_{i\in\mathcal{A}}, V'\rangle$ with

- $W' = W \cap \{w \mid \mathcal{M}, w \models \psi\}$
- For each i, $\sim'_i = \sim_i \cap (W' \times W')$
- For each i, $\preceq'_i = \preceq_i \cap (W' \times W')$
- for all $p \in \text{At}$, $V'(p) = V(p) \cap W'$

Public Announcement Logic

$$[!\psi]p \quad \leftrightarrow \quad (\psi \to p)$$

Public Announcement Logic

$$\begin{aligned} [!\psi]p &\leftrightarrow (\psi \rightarrow p) \\ [!\psi]\neg\varphi &\leftrightarrow (\psi \rightarrow \neg[!\psi]\varphi) \end{aligned}$$

Public Announcement Logic

$$\begin{array}{rcl} [!\psi]p & \leftrightarrow & (\psi \rightarrow p) \\ [!\psi]\neg\varphi & \leftrightarrow & (\psi \rightarrow \neg[!\psi]\varphi) \\ [!\psi](\varphi \wedge \chi) & \leftrightarrow & ([!\psi]\varphi \wedge [!\psi]\chi) \end{array}$$

Public Announcement Logic

$$
\begin{array}{rcl}
[!\psi]p & \leftrightarrow & (\psi \rightarrow p) \\
[!\psi]\neg\varphi & \leftrightarrow & (\psi \rightarrow \neg[!\psi]\varphi) \\
[!\psi](\varphi \wedge \chi) & \leftrightarrow & ([!\psi]\varphi \wedge [!\psi]\chi) \\
[!\psi][!\varphi]\chi & \leftrightarrow & [!(\psi \wedge [!\psi]\varphi)]\chi
\end{array}
$$

Public Announcement Logic

$$
\begin{array}{rcl}
[!\psi]p & \leftrightarrow & (\psi \rightarrow p) \\
[!\psi]\neg\varphi & \leftrightarrow & (\psi \rightarrow \neg[!\psi]\varphi) \\
[!\psi](\varphi \wedge \chi) & \leftrightarrow & ([!\psi]\varphi \wedge [!\psi]\chi) \\
[!\psi][!\varphi]\chi & \leftrightarrow & [!(\psi \wedge [!\psi]\varphi)]\chi \\
[!\psi]K_i\varphi & \leftrightarrow & (\psi \rightarrow K_i(\psi \rightarrow [!\psi]\varphi))
\end{array}
$$

Public Announcement Logic

$$
\begin{array}{rcl}
[!\psi]p & \leftrightarrow & (\psi \rightarrow p) \\
[!\psi]\neg\varphi & \leftrightarrow & (\psi \rightarrow \neg[!\psi]\varphi) \\
[!\psi](\varphi \wedge \chi) & \leftrightarrow & ([!\psi]\varphi \wedge [!\psi]\chi) \\
[!\psi][!\varphi]\chi & \leftrightarrow & [!(\psi \wedge [!\psi]\varphi)]\chi \\
[!\psi]K_i\varphi & \leftrightarrow & (\psi \rightarrow K_i(\psi \rightarrow [!\psi]\varphi))
\end{array}
$$

Theorem Every formula of Public Announcement Logic is equivalent to a formula of Epistemic Logic.

Public Announcement vs. Conditional Belief

Are $[!\varphi]B\psi$ and $B^{\varphi}\psi$ different?

Public Announcement vs. Conditional Belief

Are $[!\varphi]B\psi$ and $B^{\varphi}\psi$ different? Yes!

Public Announcement vs. Conditional Belief

Are $[!\varphi]B\psi$ and $B^{\varphi}\psi$ different? Yes!

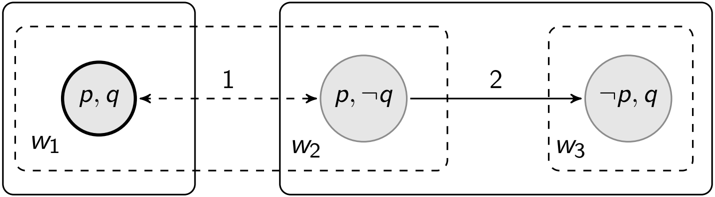

Public Announcement vs. Conditional Belief

Are $[!\varphi]B\psi$ and $B^{\varphi}\psi$ different? Yes!

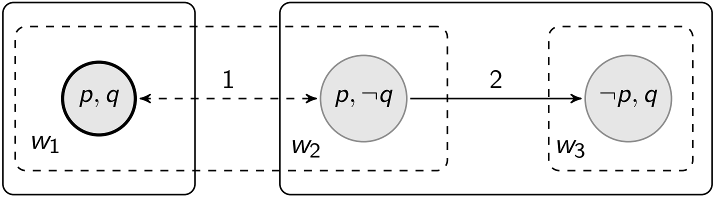

- $w_1 \models B_1 B_2 q$

Public Announcement vs. Conditional Belief

Are $[!\varphi]B\psi$ and $B^{\varphi}\psi$ different? Yes!

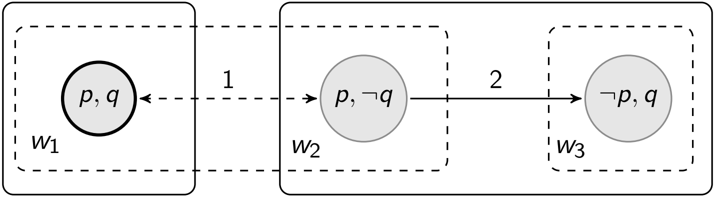

- $w_1 \models B_1 B_2 q$
- $w_1 \models B_1^p B_2 q$

Public Announcement vs. Conditional Belief

Are $[!\varphi]B\psi$ and $B^{\varphi}\psi$ different? Yes!

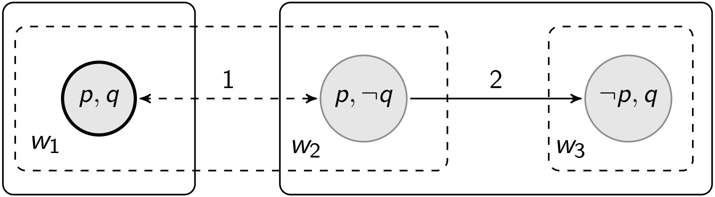

- $w_1 \models B_1 B_2 q$
- $w_1 \models B_1^p B_2 q$
- $w_1 \models [!p]\neg B_1 B_2 q$

Public Announcement vs. Conditional Belief

Are $[!\varphi]B\psi$ and $B^{\varphi}\psi$ different? Yes!

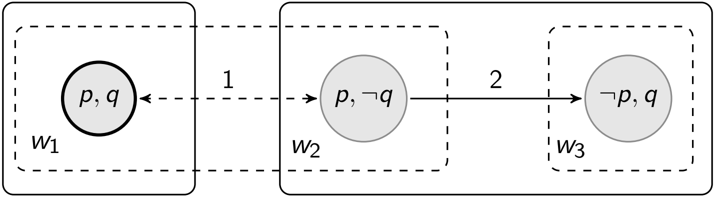

- $w_1 \models B_1 B_2 q$
- $w_1 \models B_1^p B_2 q$
- $w_1 \models [!p]\neg B_1 B_2 q$
- More generally, $B_i^p(p \wedge \neg K_i p)$ is satisfiable but $[!p]B_i(p \wedge \neg K_i p)$ is not.

The Logic of Public Observation

- $[!\varphi]K\psi \leftrightarrow (\varphi \rightarrow K(\varphi \rightarrow [!\varphi]\psi))$

The Logic of Public Observation

- $[!\varphi]K\psi \leftrightarrow (\varphi \to K(\varphi \to [!\varphi]\psi))$

- $[!\varphi][\preceq]\psi \leftrightarrow (\varphi \to [\preceq](\varphi \to [!\varphi]\psi))$

The Logic of Public Observation

- $[!\varphi]K\psi \leftrightarrow (\varphi \rightarrow K(\varphi \rightarrow [!\varphi]\psi))$

- $[!\varphi][\preceq]\psi \leftrightarrow (\varphi \rightarrow [\preceq](\varphi \rightarrow [!\varphi]\psi))$

- **Belief**: $[!\varphi]B\psi \not\leftrightarrow (\varphi \rightarrow B(\varphi \rightarrow [!\varphi]\psi))$

The Logic of Public Observation

- $[!\varphi]K\psi \leftrightarrow (\varphi \rightarrow K(\varphi \rightarrow [!\varphi]\psi))$

- $[!\varphi][\preceq]\psi \leftrightarrow (\varphi \rightarrow [\preceq](\varphi \rightarrow [!\varphi]\psi))$

- **Belief**: $[!\varphi]B\psi \not\leftrightarrow (\varphi \rightarrow B(\varphi \rightarrow [!\varphi]\psi))$

  $[!\varphi]B\psi \leftrightarrow (\varphi \rightarrow B^{\varphi}[!\varphi]\psi)$

The Logic of Public Observation

- $[!\varphi]K\psi \leftrightarrow (\varphi \rightarrow K(\varphi \rightarrow [!\varphi]\psi))$

- $[!\varphi][\preceq]\psi \leftrightarrow (\varphi \rightarrow [\preceq](\varphi \rightarrow [!\varphi]\psi))$

- **Belief**: $[!\varphi]B\psi \not\leftrightarrow (\varphi \rightarrow B(\varphi \rightarrow [!\varphi]\psi))$

  $[!\varphi]B\psi \leftrightarrow (\varphi \rightarrow B^{\varphi}[!\varphi]\psi)$
  $[!\varphi]B^{\alpha}\psi \leftrightarrow (\varphi \rightarrow B^{\varphi\wedge[!\varphi]\alpha}[!\varphi]\psi)$

Belief Revision via Plausibility

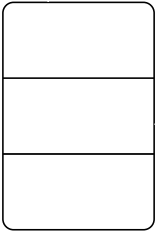

Belief Revision via Plausibility

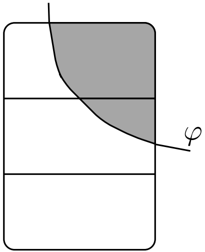

Incorporate the new information φ

Belief Revision via Plausibility

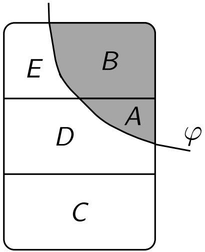

Incorporate the new information φ

Belief Revision via Plausibility

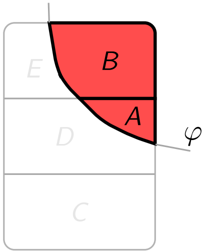

Public Announcement: Information from an infallible source

$(!\varphi)$: $A \prec_i B$

Belief Revision via Plausibility

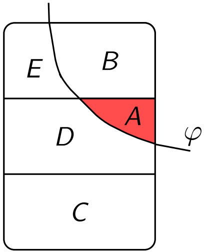

Public Announcement: Information from an infallible source
$(!\varphi)$: $A \prec_i B$

Conservative Upgrade: Information from a trusted source
$(\uparrow\varphi)$: $A \prec_i C \prec_i D \prec_i B \cup E$

Belief Revision via Plausibility

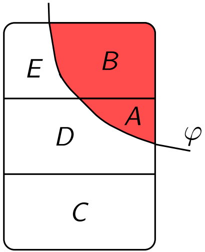

Public Announcement: Information from an infallible source
$(!\varphi)$: $A \prec_i B$

Conservative Upgrade: Information from a trusted source
$(\uparrow\varphi)$: $A \prec_i C \prec_i D \prec_i B \cup E$

Radical Upgrade: Information from a strongly trusted source
$(\Uparrow\varphi)$: $A \prec_i B \prec_i C \prec_i D \prec_i E$

$$K_0 \qquad \Longrightarrow \qquad K_t$$

$$K_0 \quad \xRightarrow{\text{Find out that } \varphi} \quad K_t = K_0 * \varphi$$

$$K_0 \quad \xRightarrow{\text{Find out that } \varphi} \quad K_t = K_0 * \varphi$$

$$\mathcal{M}_0 \quad \xRightarrow{\text{Find out that } \varphi} \quad \mathcal{M}_t = \mathcal{M}_0^{\Uparrow\varphi}$$

$$K_0 \quad \overset{\text{Find out that } \varphi}{\Longrightarrow} \quad K_t = K_0 * \varphi$$

$$\mathcal{M}_0 \quad \overset{\text{Find out that } \varphi}{\Longrightarrow} \quad \begin{aligned} \mathcal{M}_t &= \mathcal{M}_0^{\Uparrow\varphi} \\ &= \mathcal{M}_0^{\uparrow\varphi} \end{aligned}$$

$$
\begin{tikzcd}
K_0 \arrow[r, Rightarrow, "\text{Find out that } \varphi"] \arrow[d, leftrightarrow] & K_t = K_0 * \varphi \arrow[d, leftrightarrow] \\
\mathcal{M}_0 \arrow[r, Rightarrow, "\text{Find out that } \varphi"] & \mathcal{M}_t = \mathcal{M}_0^{\Uparrow\varphi} = \mathcal{M}_0^{\uparrow\varphi}
\end{tikzcd}
$$

Bayesian Models

Conceptions of Belief

Binary: "all-out" belief. For any statement p, the agent either does or does not believe p. It is natural to take an *unqualified* assertion as a statement of belief of the speaker.

Conceptions of Belief

Binary: "all-out" belief. For any statement *p*, the agent either does or does not believe *p*. It is natural to take an *unqualified* assertion as a statement of belief of the speaker.

Graded: beliefs come in degrees. We are *more confident* in some of our beliefs than in others.

Conceptions of Belief

Binary: "all-out" belief. For any statement p, the agent either does or does not believe p. It is natural to take an *unqualified* assertion as a statement of belief of the speaker.

Graded: beliefs come in degrees. We are *more confident* in some of our beliefs than in others.

Eric Schwitzgebel. *Belief*. In The Stanford Encyclopedia of Philosophy.

Franz Huber. *Formal Theories of Belief*. In The Stanford Encyclopedia of Philosophy.

Conceptions of Beliefs: Questions

What are the *formal constraints* on rational belief?

- rational graded beliefs should obey the laws of probability

Conceptions of Beliefs: Questions

What are the *formal constraints* on rational belief?

- rational graded beliefs should obey the laws of probability
- rational all-out beliefs should be consistent/deductively closed

Conceptions of Beliefs: Questions

What are the *formal constraints* on rational belief?

- rational graded beliefs should obey the laws of probability
- rational all-out beliefs should be consistent/deductively closed
- how should we justify these constraints?

D. Christensen. *Putting Logic in its Place*. Oxford University Press.

Suppose that W is a set of states (the *set of outcomes*).

A σ-algebra is a set $\Sigma \subseteq \wp(W)$ such that

- $W \in \Sigma$
- If $A \in \Sigma$, then $\overline{A} \in \Sigma$
- If $\{A_i\}$ is a countable collection of sets from Σ, then $\bigcup_i A_i \in \Sigma$

A **probability function** is a function $p : \Sigma \to [0, 1]$ satisfying:

- $p(W) = 1$
- $p(A \cup B) = p(A) + p(B)$ whenever $A \cap B = \emptyset$

(W, Σ, p) is called a probability space.

Probability

Kolmogorov Axioms:

1. For each E, $0 \leq p(E) \leq 1$
2. $p(W) = 1$, $p(\emptyset) = 0$
3. If $E_1, \ldots, E_n, \ldots$ are pairwise disjoint ($E_i \cap E_j = \emptyset$ for $i \neq j$), then $p(\bigcup_i E_i) = \sum_i p(E_i)$

Probability

Kolmogorov Axioms:

1. For each E, $0 \leq p(E) \leq 1$
2. $p(W) = 1$, $p(\emptyset) = 0$
3. If $E_1, \ldots, E_n, \ldots$ are pairwise disjoint ($E_i \cap E_j = \emptyset$ for $i \neq j$), then $p(\bigcup_i E_i) = \sum_i p(E_i)$

- $p(\overline{E}) = 1 - p(E)$ ($\overline{E}$ is the complement of E)
- If $E \subseteq F$ then $p(E) \leq p(F)$
- $p(E \cup F) = p(E) + p(F) - p(E \cap F)$

Suppose that $(\mathcal{L}, \models)$ is a logic. A probability function is a map $p : \mathcal{L} \to [0, 1]$ such that

1. For each E, $0 \leq p(\varphi) \leq 1$
2. $p(\varphi) = 1$ if $\models \varphi$
3. If $p(\varphi \vee \psi) = p(\varphi) + p(\psi)$ when $\models \neg(\varphi \wedge \psi)$.

I.J. Good. *46,656 Varieties of Bayesians*. Good Thinking: The Foundations of Probability and Its Applications, University of Minnesota Press (1983).

Conditional Probability

The probability of E *given* F, dented $p(E|F)$, is defined to be

$$p(E|F) = \frac{p(E \cap F)}{p(F)}.$$

provided $P(F) > 0$.

Bayes Theorem

$$p(H|E) = p(E|H)\frac{p(H)}{p(E)}$$

Bayes Theorem

$$p(H|E) = p(E|H)\frac{p(H)}{p(E)}$$

Bayes theorem is important because it expresses the quantity $p(H|E)$ (the probability of a hypothesis H given the evidence E) —which is something people often find hard to assess—in terms of quantities that can be drawn directly from experiential knowledge.

Example: Suppose you are in a casino and you hear a person at the next gambling table announce “Twelve”. We want to know wether he was rolling a pair of dice or a roulette wheel.

Example: Suppose you are in a casino and you hear a person at the next gambling table announce "Twelve". We want to know wether he was rolling a pair of dice or a roulette wheel.

That is, compare $p(\mathit{Dice} \mid \mathit{Twelve})$ with $p(\mathit{Roulette} \mid \mathit{Twelve})$.

Example: Suppose you are in a casino and you hear a person at the next gambling table announce "Twelve". We want to know wether he was rolling a pair of dice or a roulette wheel.

That is, compare $p(Dice \mid Twelve)$ with $p(Roulette \mid Twelve)$.

Based on our background knowledge of gambling we have $p(Twelve \mid Dice) = 1/36$ and $p(Twelve \mid Roulette) = 1/38$.

Example: Suppose you are in a casino and you hear a person at the next gambling table announce "Twelve". We want to know wether he was rolling a pair of dice or a roulette wheel.

That is, compare $p(\mathit{Dice} \mid \mathit{Twelve})$ with $p(\mathit{Roulette} \mid \mathit{Twelve})$.

Based on our background knowledge of gambling we have $p(\mathit{Twelve} \mid \mathit{Dice}) = 1/36$ and $p(\mathit{Twelve} \mid \mathit{Roulette}) = 1/38$.

Based on our observations about the casino, we can judge the prior probabilities $p(\mathit{Dice})$ and $p(\mathit{Roulette})$.

Example: Suppose you are in a casino and you hear a person at the next gambling table announce "Twelve". We want to know wether he was rolling a pair of dice or a roulette wheel.

That is, compare $p(\mathit{Dice} \mid \mathit{Twelve})$ with $p(\mathit{Roulette} \mid \mathit{Twelve})$.

Based on our background knowledge of gambling we have $p(\mathit{Twelve} \mid \mathit{Dice}) = 1/36$ and $p(\mathit{Twelve} \mid \mathit{Roulette}) = 1/38$.

Based on our observations about the casino, we can judge the prior probabilities $p(\mathit{Dice})$ and $p(\mathit{Roulette})$.

But this is now enough to *calculate* the required probabilities.

Extensions and variations

- Dempster-Shafer belief functions: $Bel : A \to [0,1]$ are *super-additive*, $Bel(A) + Bel(B) \leq Bel(A \cup B)$ if $A \cap B = \emptyset$. The the number $Bel(A)$ represents the strength with which A is supported by the agent's knowledge or belief base.

- Dempster-Shafer belief functions: $\mathit{Bel} : \mathcal{A} \to [0,1]$ are *super-additive*, $\mathit{Bel}(A) + \mathit{Bel}(B) \leq \mathit{Bel}(A \cup B)$ if $A \cap B = \emptyset$. The the number $\mathit{Bel}(A)$ represents the strength with which A is supported by the agent's knowledge or belief base.

- Non-standard probability: $\mu : \Sigma \to \mathbb{R}^*$

- Dempster-Shafer belief functions: $Bel : A \to [0, 1]$ are *super-additive*, $Bel(A) + Bel(B) \leq Bel(A \cup B)$ if $A \cap B = \emptyset$. The the number $Bel(A)$ represents the strength with which A is supported by the agent's knowledge or belief base.

- Non-standard probability: $\mu : \Sigma \to \mathbb{R}^*$

- Halpern Plausibility Functions: $\mu : \Sigma \to (D, \preceq)$.

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 $p(A) = x$.

Ambiguation If Π is a set of probability functions, the *ambiguation* of Π is the indeterminate probability function that assigns to each A

$$\sigma(A) = \{x \mid p(A) = x \text{ for some } p \in \Pi\}$$

Ambiguation If Π is a set of probability functions, the *ambiguation* of Π is the indeterminate probability function that assigns to each A

$$\sigma(A) = \{x \mid p(A) = x \text{ for some } p \in \Pi\}$$

Observation. The map that takes and indeterminate probability function to the class of its precisfications is clearly 1-1. However, the ambiguation of a set of probability functions can have precisfications not in the ambiguated set.

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.

Upper and Lower Probabilities

If σ is an indeterminate probability function, define

- **Lower probability**: $\sigma_*(A) = \inf\{x \mid x \in \sigma(A)\}$
- **Upper probability**: $\sigma^*(A) = \sup\{x \mid x \in \sigma(A)\}$

Dutch Book Arguments

Should a rational agent's graded beliefs satisfy the laws of probability?

Ramsey, de Finetti and Savage (1)

How do we *measure* a (rational) agent's subjective probabilities?

Ramsey, de Finetti and Savage (1)

How do we *measure* a (rational) agent's subjective probabilities?

Suppose we are wondering about Ann's degree of belief about whether a coin will land heads (H) or tails (T).

Ramsey, de Finetti and Savage (1)

How do we *measure* a (rational) agent's subjective probabilities?

Suppose we are wondering about Ann's degree of belief about whether a coin will land heads (H) or tails (T).

Why don't we just *ask* her?

Ramsey, de Finetti and Savage (1)

How do we *measure* a (rational) agent's subjective probabilities?

Suppose we are wondering about Ann's degree of belief about whether a coin will land heads (H) or tails (T).

Why don't we just *ask* her? reported vs. "actual" degrees of belief.

Ramsey, de Finetti and Savage (1)

How do we *measure* a (rational) agent's subjective probabilities?

Suppose we are wondering about Ann's degree of belief about whether a coin will land heads (H) or tails (T).

Why don't we just *ask* her? reported vs. "actual" degrees of belief.

What we need: systematic procedures for linking the probability calculus (graded beliefs) to claims about objectively observable behavior, such as preferences revealed by choice behavior.

Ramsey, de Finetti and Savage (2)

Suppose we are wondering about Ann's degree of belief about whether a coin will land heads (H) or tails (T).

Offer Ann two bets:

- L_1 If the coin lands heads, you win a sports car; otherwise you win nothing
- L_2 If the coin does not land heads, you win a sports car; otherwise you win nothing.

Ramsey, de Finetti and Savage (2)

Suppose we are wondering about Ann's degree of belief about whether a coin will land heads (H) or tails (T).

Offer Ann two bets:

L_1 If the coin lands heads, you win a sports car; otherwise you win nothing

L_2 If the coin does not land heads, you win a sports car; otherwise you win nothing.

If Ann chooses L_1, she believes H is more probable than T
If Ann chooses L_2, she believes T is more probable than H
If Ann is indifferent, she believes H and T are equally probable (i.e., $p_A(H) = p_A(T) = 1/2$)

The Dutch Book Argument

But, why *should* a rational agent's graded beliefs satisfy the Kolmogorov axioms?

The Dutch Book Argument

But, why *should* a rational agent's graded beliefs satisfy the Kolmogorov axioms?

Anyone whose beliefs violate the laws of probability is *practically irrational*.

F. P. Ramsey. *Truth and Probability*. 1931.

B. de Finetti. *La prévision: Ses lois logiques, ses sources subjectives*. 1937.

Alan Hájek. *Dutch Book Arguments*. Oxford Handbook of Rational and Social Choice, 2008.

The EU-Thesis

Expected Money/Value/Utility: Given an agent's beliefs and desires, the **expected utility** of an **action** leading to a set of outcomes *Out* is:

$$\sum_{o\in Out} [\text{how likely the act will lead to } o]\times[\text{how much the agent desires } o]$$

The EU-Thesis

Expected Money/Value/Utility: Given an agent's beliefs and desires, the **expected utility** of an **action** leading to a set of outcomes Out is:

$$\sum_{o \in Out} [\text{how likely the act will lead to } o] \times [\text{how much the agent desires } o]$$

1. principle of maximizing expected monetary value
2. principle of maximizing expected value
3. principle of maximizing expected utility

Betting Behavior

The *EU*-thesis entails that a person satisfying 1-3 will reveal the strengths of her beliefs in her betting behavior.

Betting Behavior

The *EU*-thesis entails that a person satisfying 1-3 will reveal the strengths of her beliefs in her betting behavior.

A **wager**: $W_X = [\ a \text{ if } X, b \text{ otherwise}]$: "you get a EUR if X is true and b EUR otherwise.
(X's truth does not depend causally on W)

Betting Behavior

The *EU*-thesis entails that a person satisfying 1-3 will reveal the strengths of her beliefs in her betting behavior.

A **wager**: $W_X = [\ a \text{ if } X, \ b \text{ otherwise}]$: "you get a EUR if X is true and b EUR otherwise.
(X's truth does not depend causally on W)

The EU-thesis entails that the agent's level of confidence in X will be revealed by the monetary value she puts on W_X.

Betting Behavior

fair price f **for** W_X: the sum of money at which she is indifferent between receiving a payment of f EUR or having W_X go into effect.

Betting Behavior

fair price f **for** W_X: the sum of money at which she is indifferent between receiving a payment of f EUR or having W_X go into effect.

$f = ExpVal(W_X) = C(X) \cdot a + (1 - C(X)) \cdot b$ implies $C(X) = \frac{f-b}{a-b}$

Betting Behavior

fair price f **for** W_X: the sum of money at which she is indifferent between receiving a payment of f EUR or having W_X go into effect.

$$f = ExpVal(W_X) = C(X) \cdot a + (1 - C(X)) \cdot b \text{ implies } C(X) = \frac{f-b}{a-b}$$

If she is indifferent between $63,81$ EUR and [100 EUR if it rains, 0 EUR otherwise], then she believes to degree 0.6381 that it will rain.

Dutch Book

An agent will swap an (set of) wagers with the (sum of) their fair prices.

Dutch Book

Suppose that X and Y are logically incompatible ($X \cap Y = \emptyset$)

Dutch Book

Suppose that X and Y are logically incompatible ($X \cap Y = \emptyset$)

$$W_X = [1 \text{ if } X, \text{ 0 else}]$$
$$W_Y = [1 \text{ if } Y, \text{ 0 else}]$$
$$W_{X \vee Y} = [1 \text{ if } X \vee Y, \text{ 0 else}]$$

Dutch Book

Suppose that X and Y are logically incompatible ($X \cap Y = \emptyset$)

$$W_X = [1 \text{ if } X, 0 \text{ else}]$$
$$W_Y = [1 \text{ if } Y, 0 \text{ else}]$$
$$W_{X \vee Y} = [1 \text{ if } X \vee Y, 0 \text{ else}]$$

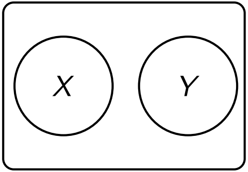

Dutch Book

Suppose that X and Y are logically incompatible ($X \cap Y = \emptyset$)

Fair price: $f = 0.25$ for $W_X = [1 \text{ if } X, 0 \text{ else}]$
Fair price: $f = 0.25$ for $W_Y = [1 \text{ if } Y, 0 \text{ else}]$
Fair price: $f = 0.6$ for $W_{X \vee Y} = [1 \text{ if } X \vee Y, 0 \text{ else}]$

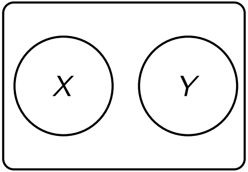

Dutch Book

Suppose that X and Y are logically incompatible ($X \cap Y = \emptyset$)

Fair price: $f = 0.25$ for $W_X = [1 \text{ if } X, 0 \text{ else}]$
Fair price: $f = 0.25$ for $W_Y = [1 \text{ if } Y, 0 \text{ else}]$
Fair price: $f = 0.6$ for $W_{X \vee Y} = [1 \text{ if } X \vee Y, 0 \text{ else}]$

Consider $\mathcal{W}_1 = \{0.6, W_X, W_Y\}$ and $\mathcal{W}_2 = \{0.5, W_{X \vee Y}\}$

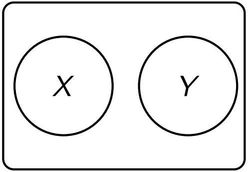

Dutch Book

Suppose that X and Y are logically incompatible ($X \cap Y = \emptyset$)

Fair price: $f = 0.25$ for $W_X = [1 \text{ if } X, 0 \text{ else}]$
Fair price: $f = 0.25$ for $W_Y = [1 \text{ if } Y, 0 \text{ else}]$
Fair price: $f = 0.6$ for $W_{X \vee Y} = [1 \text{ if } X \vee Y, 0 \text{ else}]$

Consider $\mathcal{W}_1 = \{0.6, W_X, W_Y\}$ and $\mathcal{W}_2 = \{0.5, W_{X \vee Y}\}$

- indifferent between $\mathcal{W}_1$ and $\mathcal{W}_2$

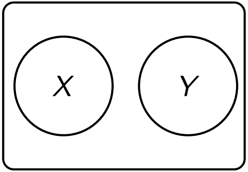

Dutch Book

Suppose that X and Y are logically incompatible ($X \cap Y = \emptyset$)

Fair price: $f = 0.25$ for $W_X = [1 \text{ if } X,\ 0 \text{ else}]$
Fair price: $f = 0.25$ for $W_Y = [1 \text{ if } Y,\ 0 \text{ else}]$
Fair price: $f = 0.6$ for $W_{X \vee Y} = [1 \text{ if } X \vee Y,\ 0 \text{ else}]$

Consider $\mathcal{W}_1 = \{0.6, W_X, W_Y\}$ and $\mathcal{W}_2 = \{0.5, W_{X \vee Y}\}$

- indifferent between $\mathcal{W}_1$ and $\mathcal{W}_2$
- swap $\mathcal{W}_1$ for $\mathcal{W}_2$

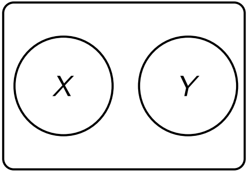

Dutch Book

Suppose that X and Y are logically incompatible ($X \cap Y = \emptyset$)

Fair price: $f = 0.25$ for $W_X = [1 \text{ if } X, 0 \text{ else}]$
Fair price: $f = 0.25$ for $W_Y = [1 \text{ if } Y, 0 \text{ else}]$
Fair price: $f = 0.6$ for $W_{X \vee Y} = [1 \text{ if } X \vee Y, 0 \text{ else}]$

Consider $\mathcal{W}_1 = \{0.6, W_X, W_Y\}$ and $\mathcal{W}_2 = \{0.5, W_{X\vee Y}\}$

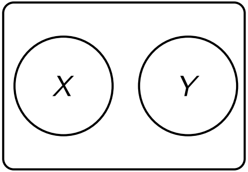

- indifferent between $\mathcal{W}_1$ and $\mathcal{W}_2$
- swap $\mathcal{W}_1$ for $\mathcal{W}_2$
- But $\mathcal{W}_1$ is always better:

Dutch Book

Suppose that X and Y are logically incompatible ($X \cap Y = \emptyset$)

Fair price: $f = 0.25$ for $W_X = [1 \text{ if } X, 0 \text{ else}]$
Fair price: $f = 0.25$ for $W_Y = [1 \text{ if } Y, 0 \text{ else}]$
Fair price: $f = 0.6$ for $W_{X \vee Y} = [1 \text{ if } X \vee Y, 0 \text{ else}]$

Consider $\mathcal{W}_1 = \{0.6, W_X, W_Y\}$ and $\mathcal{W}_2 = \{0.5, W_{X \vee Y}\}$

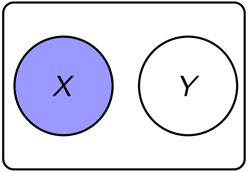

- indifferent between $\mathcal{W}_1$ and $\mathcal{W}_2$
- swap $\mathcal{W}_1$ for $\mathcal{W}_2$
- But $\mathcal{W}_1$ is always better:
  - If X is true
    payoff($\mathcal{W}_1$) $= 1.6 >$ payoff($\mathcal{W}_2$)$=1.5$

Dutch Book

Suppose that X and Y are logically incompatible ($X \cap Y = \emptyset$)

Fair price: $f = 0.25$ for $W_X = [1 \text{ if } X, 0 \text{ else}]$
Fair price: $f = 0.25$ for $W_Y = [1 \text{ if } Y, 0 \text{ else}]$
Fair price: $f = 0.6$ for $W_{X \vee Y} = [1 \text{ if } X \vee Y, 0 \text{ else}]$

Consider $\mathcal{W}_1 = \{0.6, W_X, W_Y\}$ and $\mathcal{W}_2 = \{0.5, W_{X \vee Y}\}$

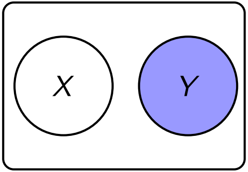

- indifferent between $\mathcal{W}_1$ and $\mathcal{W}_2$
- swap $\mathcal{W}_1$ for $\mathcal{W}_2$
- But $\mathcal{W}_1$ is always better:
  - If X is true
    payoff($\mathcal{W}_1$) $= 1.6 >$ payoff($\mathcal{W}_2$)$=1.5$
  - If Y is true
    payoff($\mathcal{W}_1$) $= 1.6 >$ payoff($\mathcal{W}_2$)$= 1.5$

Dutch Book

Suppose that X and Y are logically incompatible ($X \cap Y = \emptyset$)

Fair price: $f = 0.25$ for $W_X = [1 \text{ if } X, 0 \text{ else}]$
Fair price: $f = 0.25$ for $W_Y = [\ 1 \text{ if } Y, 0 \text{ else}]$
Fair price: $f = 0.6$ for $W_{X \vee Y} = [1 \text{ if } X \vee Y, 0 \text{ else}]$

Consider $\mathcal{W}_1 = \{0.6, W_X, W_Y\}$ and $\mathcal{W}_2 = \{0.5, W_{X \vee Y}\}$

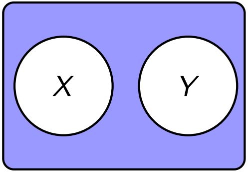

- indifferent between $\mathcal{W}_1$ and $\mathcal{W}_2$
- swap $\mathcal{W}_1$ for $\mathcal{W}_2$
- But $\mathcal{W}_1$ is always better:
  - If X is true
    payoff($\mathcal{W}_1$) $= 1.6 >$ payoff($\mathcal{W}_2$)$=1.5$
  - If Y is true
    payoff($\mathcal{W}_1$) $= 1.6 >$ payoff($\mathcal{W}_2$)$= 1.5$
  - If neither X nor Y is true
    payoff($\mathcal{W}_1$) $= 0.6 >$ payoff($\mathcal{W}_2$)$=0.5$

Dutch Book Theorem

Theorem. Imagine and EU-maximizer who satisfies 1-3 and has a precise degree of belief for every proposition she considers. If these beliefs violate the laws of probability, then she will make Dutch Book against herself.

Dutch Book Theorem

Theorem. Imagine and EU-maximizer who satisfies 1-3 and has a precise degree of belief for every proposition she considers. If these beliefs violate the laws of probability, then she will make Dutch Book against herself.

This assumes there is an agent who

1. meets conditions 1-3

Dutch Book Theorem

Theorem. Imagine and EU-maximizer who satisfies 1-3 and has a precise degree of belief for every proposition she considers. If these beliefs violate the laws of probability, then she will make Dutch Book against herself.

This assumes there is an agent who

1. meets conditions 1-3
2. sets a fair price for every wager she considers

Dutch Book Theorem

Theorem. Imagine and EU-maximizer who satisfies 1-3 and has a precise degree of belief for every proposition she considers. If these beliefs violate the laws of probability, then she will make Dutch Book against herself.

This assumes there is an agent who

1. meets conditions 1-3
2. sets a fair price for every wager she considers
3. maximizes expected utility

Dutch Book Theorem

Theorem. Imagine and EU-maximizer who satisfies 1-3 and has a precise degree of belief for every proposition she considers. If these beliefs violate the laws of probability, then she will make Dutch Book against herself.

This assumes there is an agent who

1. meets conditions 1-3
2. sets a fair price for every wager she considers
3. maximizes expected utility

allow agents to have incomplete or imprecise preferences

Dutch Book Theorem

Theorem. Imagine and EU-maximizer who satisfies 1-3 and has a precise degree of belief for every proposition she considers. If these beliefs violate the laws of probability, then she will make Dutch Book against herself.

This assumes there is an agent who

1. meets conditions 1-3
2. sets a fair price for every wager she considers
3. maximizes expected utility

justify probabilistic coherence and EU simultaneously: Savage's Representation Theorem

J. Joyce. *A nonpragmatic vindication of probabilism*. Philosophy of Science 65, 575603 (1998).

R. Pettigrew. *Epistemic Utility Arguments for Probabilism*. Stanford Encyclopedia of Philosophy, 2015.

R. Pettigrew. *Accuracy and the Laws of Credence*. Oxford University Press, 2016.

Accuracy

Accuracy. An epistemic agent ought to approximate the truth. In other words, she ought to minimize her inaccuracy.

Accuracy

Accuracy (Synchronic expected local). An agent ought to minimize the **expected local inaccuracy** of her degrees of credence in all propositions $A \subseteq W$ by the lights of her current belief function, relative to a **legitimate local inaccuracy measure** and over the set of worlds that are currently epistemically possible for her.

Accuracy (Synchronic expected global). An agent ought to minimize the **expected global inaccuracy** of her current belief function by the lights of her current belief function, relative to a legitimate global inaccuracy measure and over the set of worlds that are currently epistemically possible for her

Measuring Inaccuracy

Alethic Vindication The ideal credence function at world w is the omniscient credence function at w, namely, v_w.

Perfectionism The accuracy of a credence function at a world is its proximity to the ideal credence function at that world.

Squared Euclidean Distance Distance between credence functions is measured by squared Euclidean distance.

B. de Finetti. *Theory of Probability*. John Wiley and Sons, 1974.

J. Pred, R. Seiringer, E. Lieb, D. Osherson, H. V. Poor, and S. Kulkarni. *Probabilistic Coherence and Proper Scoring Rules*. IEEE Transactions on Information Theory, 2009.

Consider a vector $\mathcal{E} = (E_1, \ldots, E_n)$ of events.

A **forecast** for $\mathcal{E}$ is a vector $\mathbf{f} = (f_1, \ldots, f_n)$.

Consider a vector $\mathcal{E} = (E_1, \ldots, E_n)$ of events.

A **forecast** for $\mathcal{E}$ is a vector $\mathbf{f} = (f_1, \ldots, f_n)$.

Two possible defects:

1. There may be a rival forecast $\mathbf{g}$ that guarantees a lower **penalty** than the one for $\mathbf{f}$, regardless of which events come to pass.
2. The events in $\mathcal{E}$ may be related by inclusion or partition and $\mathbf{f}$ might violate constraints imposed by probability.

Consider a vector $\mathcal{E} = (E_1, \ldots, E_n)$ of events.

A **forecast** for $\mathcal{E}$ is a vector $\mathbf{f} = (f_1, \ldots, f_n)$.

Two possible defects:

1. There may be a rival forecast $\mathbf{g}$ that guarantees a lower **penalty** than the one for $\mathbf{f}$, regardless of which events come to pass.
2. The events in $\mathcal{E}$ may be related by inclusion or partition and $\mathbf{f}$ might violate constraints imposed by probability.

de Finetti, Predd et al., Lindley, ...: The two defects are equivalent.

Brier Score

$$\mathcal{E} = (E, F) \text{ with } E \subseteq F$$
$$\mathbf{f} = (0.6, 0.9)$$

Brier Score

$$\mathcal{E} = (E, F) \text{ with } E \subseteq F$$

$$\mathbf{f} = (0.6, 0.9)$$

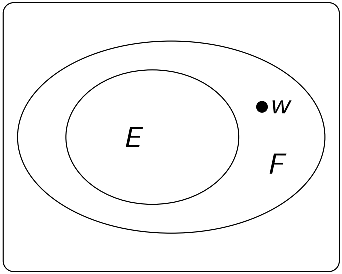

Brier Score

$$\mathcal{E} = (E, F) \text{ with } E \subseteq F$$

$$\mathbf{f} = (0.6, 0.9)$$

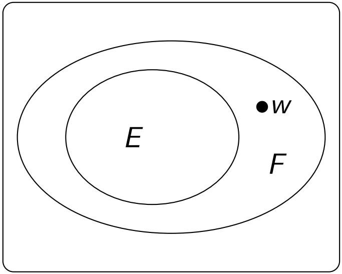

Penalty: $(0 - 0.6)^2 + (1 - 0.9)^2 = 0.37$

Proper Scoring Rule

$\mathcal{E} = (E, F)$ with $E \subseteq F$

$\mathbf{f} = (0.6, 0.9)$

Expected Penalty for E:
$0.6 * (1 - 0.6)^2 + 0.4 * (0 - 0.6)^2 = 0.230$

Expected Penalty for E **by lying**:
$0.6 * (1 - 0.65)^2 + 0.4 * (0 - 0.65)^2 = 0.2425$

Proper Scoring Rule

Suppose your probability for an event E is p, that your announced probability is x, and that your penalty assessed according ot the rule $(1-x)^2$ if E comes out true; $(0-x)^2$ otherwise. Then your expected penalty is uniquely minimized by choosing $x = p$.

Proper Scoring Rule

Suppose your probability for an event E is p, that your announced probability is x, and that your penalty assessed according ot the rule $(1-x)^2$ if E comes out true; $(0-x)^2$ otherwise. Then your expected penalty is uniquely minimized by choosing $x = p$.

<u>*Absolute Deviation*</u>

Expected Penalty for E:
$0.6 * |1 - 0.6| + 0.4 * |0 - 0.6| = 0.48$

Expected Penalty for E by lying:
$0.6 * |1 - 0.65| + 0.4 * |0 - 0.65| = 0.47$

$\mathcal{E} = (E, F)$ with $E \subseteq F$

$\mathbf{f} = (0.6, 0.9)$
$\mathbf{f}' = (0.95, 0.55)$

$\mathcal{E} = (E, F)$ with $E \subseteq F$

$\mathbf{f} = (0.6, 0.9)$
$\mathbf{f}' = (0.95, 0.55)$

Penalties:

Possibility	$\mathbf{f}$	$\mathbf{f}'$
E true, F true	0.17	0.205
E false, F true	0.37	1.105
E false, F false	1.17	1.205

S is a sample space. Subsets of S are events. Let $\mathcal{E} = (E_1, \ldots, E_n)$ be a vector of events.

S is a sample space. Subsets of S are events. Let $\mathcal{E} = (E_1, \ldots, E_n)$ be a vector of events.

A forecast is an element of $[0, 1]^n$. A forecast is **coherent** if there is a probability measure μ over S such that for all $i = 1, \ldots, n$, $f_i = \mu(E_i)$.

S is a sample space. Subsets of S are events. Let $\mathcal{E} = (E_1, \ldots, E_n)$ be a vector of events.

A forecast is an element of $[0,1]^n$. A forecast is **coherent** if there is a probability measure μ over S such that for all $i = 1, \ldots, n$, $f_i = \mu(E_i)$.

A function $s : \{0,1\} \times [0,1] \to [0,\infty]$ is a **proper scoring rule** in case:

1. $ps(1,x) + (1-p)s(0,x)$ si uniquely minimized at $x = p$ for $p \in [0,1]$.
2. s is continuous. For $i \in \{0,1\}$, $\lim_{n\to\infty} s(i,x_n) = s(i,x)$ for any sequence x_n from $[0,1]$ converging to x.

Penalty

Given a proper scoring rule s, the penalty P_s based on s for forecast $\mathbf{f}$ and $w \in S$ is given by:

$$P_s(w, \mathbf{f}) = \sum_{i=1}^{n} s(\chi_{E_i}(w), f_i)$$

$\mathbf{f}$ is **weakly dominated** by $\mathbf{g}$ in case $P_s(w, \mathbf{g}) \leq P_s(w, \mathbf{f})$ for all $w \in S$.

$\mathbf{f}$ is **weakly dominated** by $\mathbf{g}$ in case $P_s(w, \mathbf{g}) < P_s(w, \mathbf{f})$ for all $w \in S$.

Theorem Let $\mathbf{f}$ be a forecast.

1. If $\mathbf{f}$ is coherent, then it is not weakly dominated by any forecast $\mathbf{g} \neq \mathbf{f}$
2. If $\mathbf{f}$ is incoherent, then it is strongly dominated by some coherent forecast $\mathbf{g}$

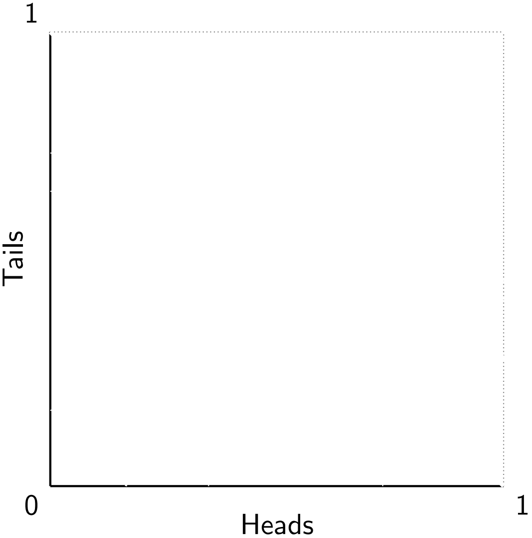

w_H: The coin is facing heads up.

w_T: The coin is facing tails up.

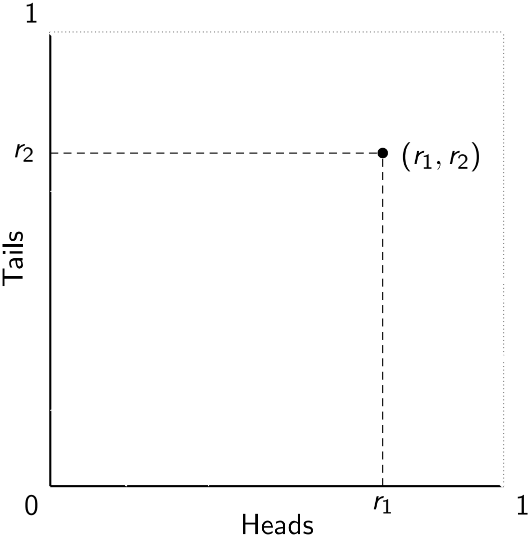

w_H: The coin is facing heads up.

w_T: The coin is facing tails up.

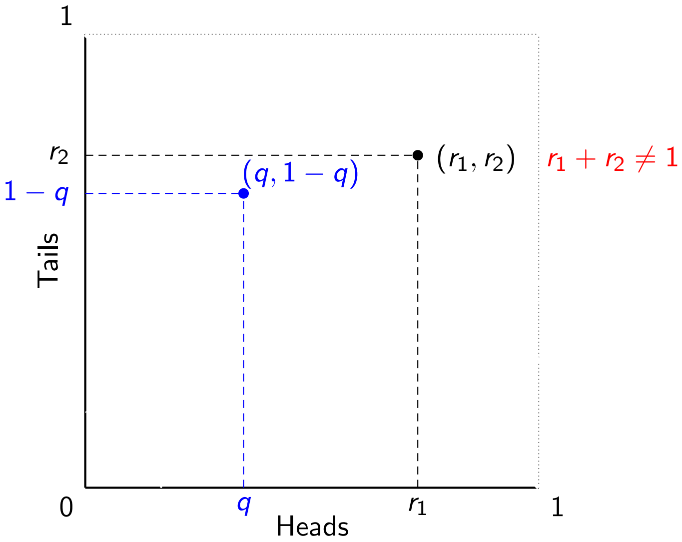

w_H: The coin is facing heads up.

w_T: The coin is facing tails up.

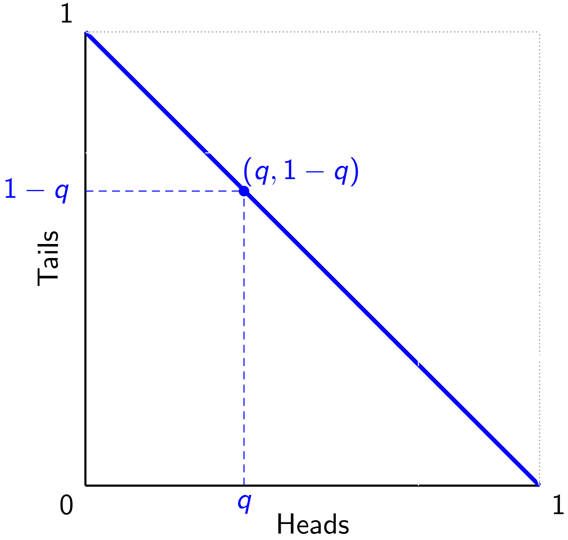

w_H: The coin is facing heads up.

w_T: The coin is facing tails up.

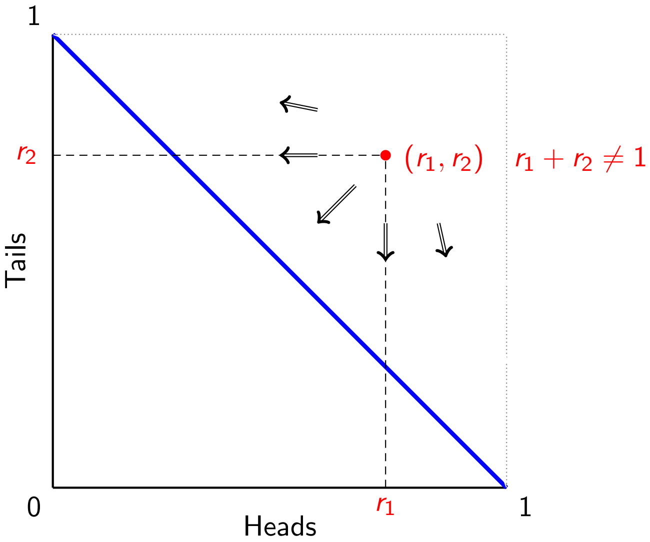

w_H: The coin is facing heads up.

w_T: The coin is facing tails up.

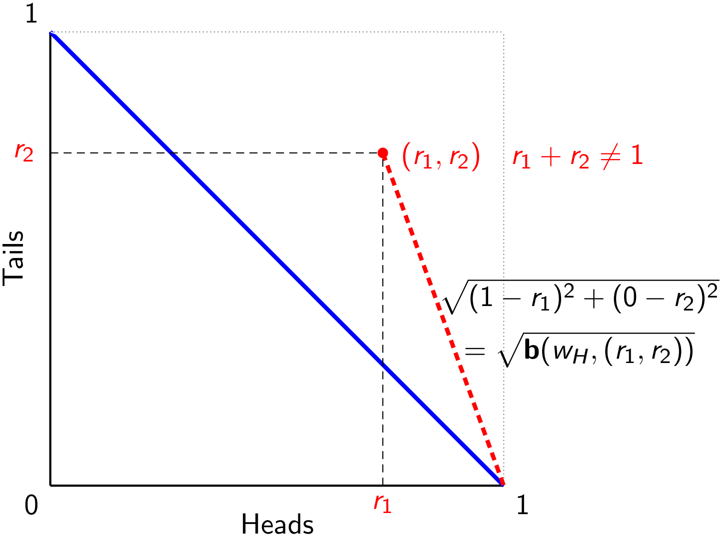

w_H: The coin is facing heads up.

w_T: The coin is facing tails up.

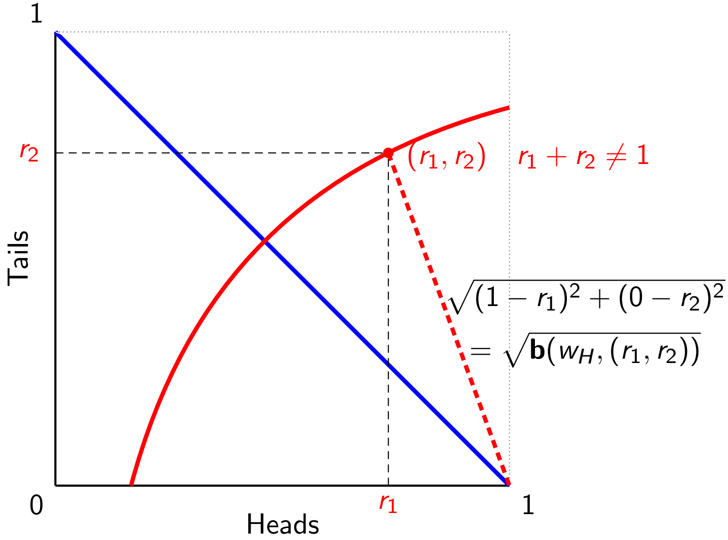

w_H: The coin is facing heads up.

w_T: The coin is facing tails up.

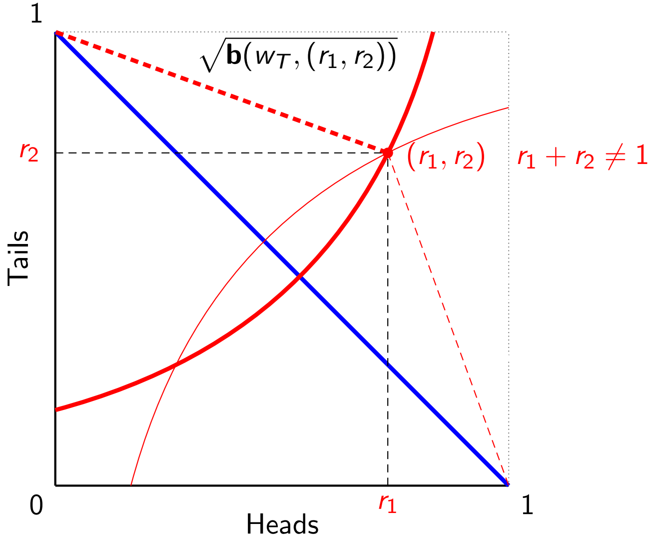

w_H: The coin is facing heads up.

w_T: The coin is facing tails up.

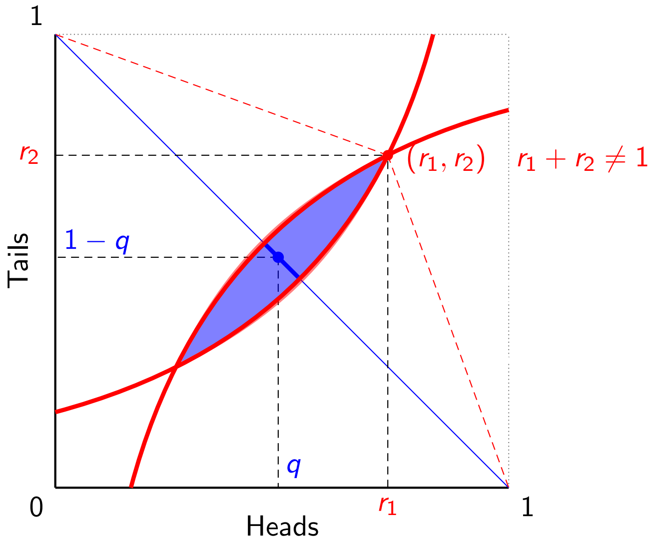

w_H: The coin is facing heads up.

w_T: The coin is facing tails up.

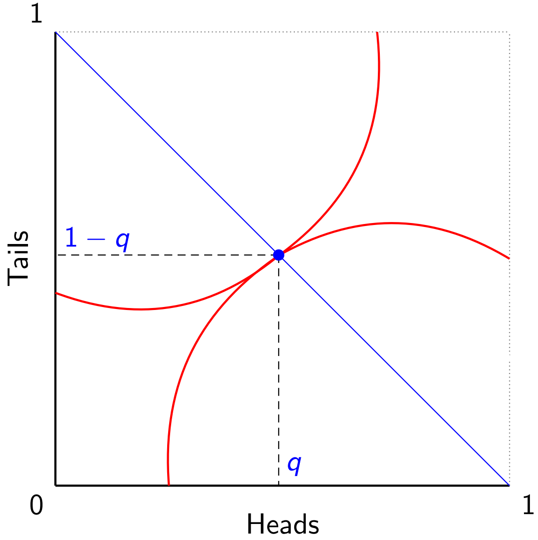

w_H: The coin is facing heads up.

w_T: The coin is facing tails up.

$$K_0 \qquad \Longrightarrow \qquad K_t$$

$$K_0 \quad \xRightarrow[\text{Suppose that } \varphi]{\text{Learn that } \varphi} \quad K_t = K_0 * \varphi$$

$$K_0 \xRightarrow[\text{}]{\text{Learn that } \varphi \text{ / Suppose that } \varphi} K_t = K_0 * \varphi$$

$$p_0 \Longrightarrow p_t$$

$$K_0 \quad \overset{\substack{\text{Learn that } \varphi \\ \text{Suppose that } \varphi}}{\Longrightarrow} \quad K_t = K_0 * \varphi$$

$$p_0 \quad \overset{\text{Learning experience}}{\Longrightarrow} \quad p_t = ???$$

$$\begin{tikzcd}[column sep=huge, row sep=huge]
K_0 \arrow[r, Rightarrow, "\begin{matrix}\text{Learn that } \varphi \\ \text{Suppose that } \varphi\end{matrix}"] \arrow[d, leftrightarrow] & K_t = K_0 * \varphi \arrow[d, leftrightarrow] \\
p_0 \arrow[r, Rightarrow, "\text{Learning experience}"] & p_t = ???
\end{tikzcd}$$

Conditioning

Conditional Probability

The probability of E *given* F, dented $p(E|F)$, is defined to be

$$p(E|F) = \frac{p(E \cap F)}{p(F)}.$$

provided $P(F) > 0$.

Conditioning

When you acquire new evidence E, the new probability of any proposition H should be the previous conditional probability of H given E. I.e., $q(H) = p(H \mid E)$.

- If p is a probability function, and $q(H) = p(H \mid E)$ for each H, then q is a probability function.

Conditioning

When you acquire new evidence E, the new probability of any proposition H should be the previous conditional probability of H given E. I.e., $q(H) = p(H \mid E)$.

- If p is a probability function, and $q(H) = p(H \mid E)$ for each H, then q is a probability function.

- (Assuming E_1 and E_2 are consistent) If q comes from p by conditioning on E_1 and r comes from q by conditioning on E_2, the result of condition on E_2 first then E_1 would have been the same, namely $r(\cdot) = p(\cdot \mid E_1 \cap E_2)$.

Setting $p_t(\cdot) = p_0(\cdot \mid E)$ is demonstrably the correct thing to do just in case, for all propositions $H \in \Sigma$, both:

1. Certainty: $p_t(E) = 1$
2. Rigidity: $p_t(H \mid E) = p_0(H \mid E)$

People are often not aware of all that they have learnt or they fail to adequately represent it, and it is only the failure of the Rigidity condition that alerts us to this.

Three Prisoner's Problem

Three prisoners A, B and C have been tried for murder and their verdicts will told to them tomorrow morning. They know only that one of them will be declared guilty and will be executed while the others will be set free. The identity of the condemned prisoner is revealed to the very reliable prison guard, but not to the prisoners themselves. Prisoner A asks the guard "Please give this letter to one of my friends — to the one who is to be released. We both know that at least one of them will be released".

Three Prisoner's Problem

An hour later, A asks the guard "Can you tell me which of my friends you gave the letter to? It should give me no clue regarding my own status because, regardless of my fate, each of my friends had an equal chance of receiving my letter." The guard told him that B received his letter.

Prisoner A then concluded that the probability that he will be released is $1/2$ (since the only people without a verdict are A and C).

Three Prisoner's Problem

But, A thinks to himself:

Three Prisoner's Problem

But, A thinks to himself:

> *Before I talked to the guard my chance of being executed was 1 in 3. Now that he told me B has been released, only C and I remain, so my chances of being executed have gone from 33.33% to 50%. What happened? I made certain not to ask for any information relevant to my own fate...*

Three Prisoner's Problem

But, A thinks to himself:

> *Before I talked to the guard my chance of being executed was 1 in 3. Now that he told me B has been released, only C and I remain, so my chances of being executed have gone from 33.33% to 50%. What happened? I made certain not to ask for any information relevant to my own fate...*

Explain what is wrong with A's reasoning.

A's reasoning

Consider the following events:

G_A: "Prisoner A will be declared guilty" (we have $p(G_A) = 1/3$)

I_B: "Prisoner B will be declared innocent" (we have $p(I_B) = 2/3$)

A's reasoning

Consider the following events:

G_A: "Prisoner A will be declared guilty" (we have $p(G_A) = 1/3$)

I_B: "Prisoner B will be declared innocent" (we have $p(I_B) = 2/3$)

We have $p(I_B \mid G_A) = 1$: "If A is declared guilty then B will be declared innocent."

A's reasoning

Consider the following events:

G_A: "Prisoner A will be declared guilty" (we have $p(G_A) = 1/3$)

I_B: "Prisoner B will be declared innocent" (we have $p(I_B) = 2/3$)

We have $p(I_B \mid G_A) = 1$: "If A is declared guilty then B will be declared innocent."

Bayes Theorem:

$$p(G_A \mid I_B) =$$

A's reasoning

Consider the following events:

G_A: "Prisoner A will be declared guilty" (we have $p(G_A) = 1/3$)

I_B: "Prisoner B will be declared innocent" (we have $p(I_B) = 2/3$)

We have $p(I_B \mid G_A) = 1$: "If A is declared guilty then B will be declared innocent."

Bayes Theorem:

$$p(G_A \mid I_B) = p(I_B \mid G_A)\frac{p(G_A)}{p(I_B)} =$$

A's reasoning

Consider the following events:

G_A: "Prisoner A will be declared guilty" (we have $p(G_A) = 1/3$)

I_B: "Prisoner B will be declared innocent" (we have $p(I_B) = 2/3$)

We have $p(I_B \mid G_A) = 1$: "If A is declared guilty then B will be declared innocent."

Bayes Theorem:

$$p(G_A \mid I_B) = p(I_B \mid G_A)\frac{p(G_A)}{p(I_B)} = 1 \cdot \frac{1/3}{2/3} = 1/2$$

A's reasoning, corrected

But, A did not receive the information that B will be declared innocent, but rather that "the guard said that B will be declared innocent." So, A should have conditioned on the event:

I'_B: "The guard said that B will be declared innocent"

A's reasoning, corrected

But, A did not receive the information that B will be declared innocent, but rather that "the guard said that B will be declared innocent." So, A should have conditioned on the event:

I'_B: "The guard said that B will be declared innocent"

Given that $p(I'_B \mid G_A)$ is $1/2$ (given that A is guilty, there is a 50-50 chance that the guard could have given the letter to B or C).

A's reasoning, corrected

But, A did not receive the information that B will be declared innocent, but rather that "the guard said that B will be declared innocent." So, A should have conditioned on the event:

I'_B: "The guard said that B will be declared innocent"

Given that $p(I'_B \mid G_A)$ is $1/2$ (given that A is guilty, there is a 50-50 chance that the guard could have given the letter to B or C). This gives us the following correct calculation:

$$p(G_A \mid I'_B) = p(I'_B \mid G_A)\frac{p(G_A)}{p(I'_B)} = 1/2 \cdot \frac{1/3}{1/2} = 1/3$$

Setting $p_t(\cdot) = p_0(\cdot \mid E)$ is demonstrably the correct thing to do just in case, for all propositions $H \in \Sigma$, both:

1. Certainty: $p_t(E) = 1$
2. Rigidity: $p_t(H \mid E) = p_0(H \mid E)$

Observation by candlelight

An agent inspects a piece of cloth by candlelight, and gets the impression that it is green (G), although he concedes that it might be blue (B) or even (but very improbably) violet (V).

Observation by candlelight

An agent inspects a piece of cloth by candlelight, and gets the impression that it is green (G), although he concedes that it might be blue (B) or even (but very improbably) violet (V).

$$p_0(G) = p_0(B) = 0.3,\ p_0(V) = 0.4$$

$$\Downarrow$$

$$p_t(G) = 0.7,\ p_t(B) = 0.25,\ p_t(V) = .05$$

Observation by candlelight

An agent inspects a piece of cloth by candlelight, and gets the impression that it is green (G), although he concedes that it might be blue (B) or even (but very improbably) violet (V).

$$p_0(G) = p_0(B) = 0.3,\ \ p_0(V) = 0.4$$

$$\Downarrow$$

$$p_t(G) = 0.7,\ \ p_t(B) = 0.25,\ \ p_t(V) = .05$$

Is there a proposition E such that $p_t(\cdot) = p_0(\cdot \mid E)$?

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 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 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 $q(\cdot) = p(\cdot \mid E)$.

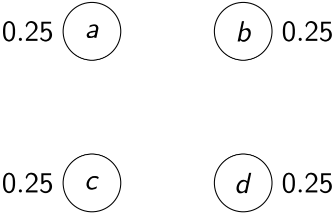
0.25
a
b
0.25
0.25
c
d
0.25

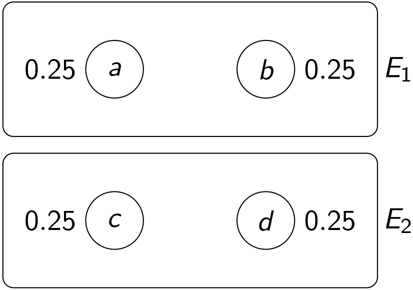

The probability that the guilty party is left-handed is 0.8

$$E_1 = \{a, b\},\ E_2 = \{c, d\}$$

$$p(E_1) = 0.8 \qquad p(E_2) = 0.2$$

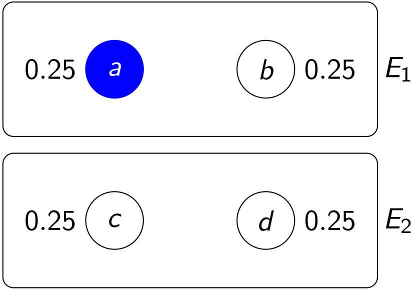

The probability that the guilty party is left-handed is 0.8

$E_1 = \{a, b\},\ E_2 = \{c, d\}$

$$p(E_1) = 0.8 \qquad p(E_2) = 0.2$$

$$p(a) = p_0(\{a\} \mid E_1) * p(E_1) + p_0(\{a\} \mid E_2) * p(E_2) = 0.5 * 0.8 + 0 = 0.4$$

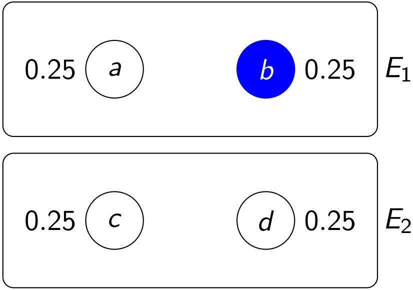

The probability that the guilty party is left-handed is 0.8

$$E_1 = \{a, b\},\ E_2 = \{c, d\}$$

$$p(E_1) = 0.8 \qquad p(E_2) = 0.2$$

$$p(b) = p_0(\{b\} \mid E_1) * p(E_1) + p_0(\{b\} \mid E_2) * p(E_2) = 0 + 0.5 * 0.8 = 0.4$$

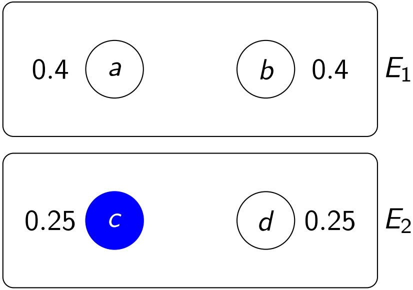

The probability that the guilty party is left-handed is 0.8

$E_1 = \{a, b\}$, $E_2 = \{c, d\}$

$$p(E_1) = 0.8 \qquad p(E_2) = 0.2$$

$$p(c) = p_0(\{c\} \mid E_1) * p(E_1) + p_0(\{c\} \mid E_2) * p(E_2) = 0 + 0.5 * 0.2 = 0.1$$

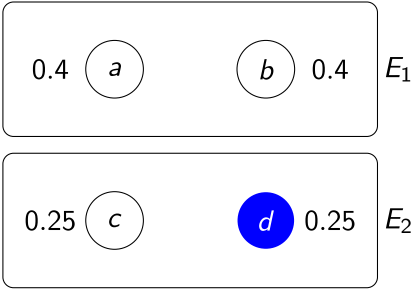

The probability that the guilty party is left-handed is 0.8

$E_1 = \{a, b\}$, $E_2 = \{c, d\}$

$p(E_1) = 0.8 \qquad p(E_2) = 0.2$

$p(d) = p_0(\{d\} \mid E_1) * p(E_1) + p_0(\{d\} \mid E_2) * p(E_2) = 0 + 0.5 * 0.2 = 0.1$

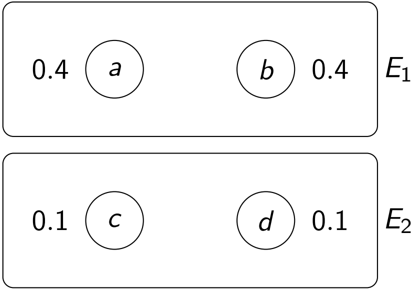

The probability that the guilty party is left-handed is 0.8

$E_1 = \{a, b\}$, $E_2 = \{c, d\}$

$$p(E_1) = 0.8 \qquad p(E_2) = 0.2$$

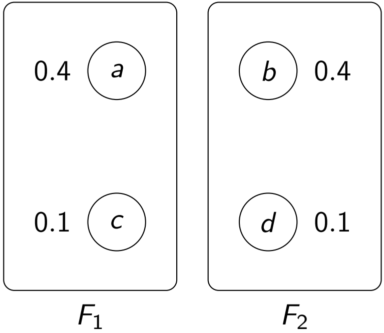

The probability that the guilty party is left-handed is 0.8

$$F_1 = \{a, c\},\ F_2 = \{b, d\}$$

$$p(F_1) = 0.7 \qquad p(F_2) = 0.3$$

$$p(a) = p_0(\{a\} \mid F_1) * p(F_1) + p_0(\{a\} \mid F_2) * p(F_2) = 0.8 * 0.7 + 0 = 0.56$$

P. Diaconis and S. Zabell. *Updating Subjective Probability*. Journal of the American Statistical Association, Vol. 77, No. 380., pp. 822-830 (1982).

Suppose we are thinking about three trials of a new surgical procedure. Under the usual circumstances a probability assignment is made on the eight possible outcomes $R = \{000, 001, 010, 011, 100, 101, 110, 111\}$, where 1 denotes a successful outcome, 0 not.

Suppose we are thinking about three trials of a new surgical procedure. Under the usual circumstances a probability assignment is made on the eight possible outcomes $R = \{000, 001, 010, 011, 100, 101, 110, 111\}$, where 1 denotes a successful outcome, 0 not.

Suppose a colleague informs us that another hospital had performed this type of operation 100 times, with 80 successful outcomes. This is clearly relevant information and we obviously want to revise our opinion.

Suppose we are thinking about three trials of a new surgical procedure. Under the usual circumstances a probability assignment is made on the eight possible outcomes $R = \{000, 001, 010, 011, 100, 101, 110, 111\}$, where 1 denotes a successful outcome, 0 not.

Suppose a colleague informs us that another hospital had performed this type of operation 100 times, with 80 successful outcomes. This is clearly relevant information and we obviously want to revise our opinion.

The information cannot be put in terms of the occurrence of an event in the original eight-point space R, and the Bayes rule is not directly available.

1.*Complete Reassessment*. In the absence of further structure it is always possible to react to the new information by completely reassessing P^*, presumably using the same techniques used to quantify the original distribution P.

2. *Retrospective Conditioning*. Some subjectivists have suggested trying to analyze this kind of problem by momentarily disregarding the new information, quantifying a distribution on a space W^* rich enough to allow ordinary conditioning to be used, and then using Bayes' rule.

3. *Exchangeability*. The three future trials may be regarded as exchangeable with the 100 trials reported by our colleague. Standard Bayesian computations can then be used. However, given that the operations will have been performed at two, possibly very different, hospitals with possibly very different patient populations, this assumption might very well be judged unsatisfactory.

4. *Jeffrey's Rule*. Suppose that the original probability assignment P was exchangeable. That is, $P(001) = P(010) = P(100)$ and $P(110) = P(101) = P(011)$. Consider a partition $\{E_i\}_{i=0}^{3}$, where $E_0 = \{000\}$, $E_1 = \{001, 010, 100\}$, $E_2 = \{110, 101, 011\}$ and $E_3 = \{111\}$. To complete the probability assignment P^*, we need a subjective assessment of each $P^*(E_i)$, then use Jeffrey's Rule to define a full probability measure.

Conditioning

- If q comes from p by conditioning on E, then $q(E) = 1$

Conditioning

- If q comes from p by conditioning on E, then $q(E) = 1$

- If $p(E_1) = 1$ then $p(E_1 \mid E_2) = 1$ for any E_2 consistent with E_1

Conditioning

- If q comes from p by conditioning on E, then $q(E) = 1$
- If $p(E_1) = 1$ then $p(E_1 \mid E_2) = 1$ for any E_2 consistent with E_1
- If $p(E_1) = 1$ then $p(A \mid E_2)$ is undefined whenever E_2 is inconsistent with E_1, since $p(E_2) = 0$

Fact. Jeffrey conditioning is not commutative.

Commutativity on Experiences Any rule for updating degrees of belief on experiences should be such that the result of updating credences on one experience and then another should be the same as the result of updating on the same two experiences in reverse order.

Holism For any experience and any proposition, there is a "defeater" proposition, such that your degree of belief in the first proposition, upon having the experience, should depend on your degree of belief in the defeater proposition.

J. Weisberg. *Commutativity or Holism? A Dilemma for Conditionalizers*. British Journal of the Philosophy of Science, 60(4), pp. 793-812, 2009.

M. Lange. *Is Jeffrey Conditionalization Defective in Virtue of Being NonCommutative? Remarks on the Sameness of Sensory Experience*. Synthese 123: 393-403, 2000.

C. Wagner. *Probability kinematics and commutativity*. Philosophy of Science 69, 266-278, 2002.

Updating probabilities

Orthodox Bayesian Policy

- accept as admissible input only propositions;
- as response to such an input the only admissible change is conditioning the prior on the proposition in question.

Updating probabilities

Orthodox Bayesian Policy

- accept as admissible input only propositions;
- as response to such an input the only admissible change is conditioning the prior on the proposition in question.

Departing from a (orthodox) Bayesian policy:

1. accept as admissible a wider variety of inputs (e.g. expected values);
2. an admissible response to such an input can be a change in the prior that is not the result of conditioning;
3. an admissible response to such an input may be non-unique, that is, the posterior may not be uniquely determined by the prior + input.

$$p_0$$

$$\boxed{\begin{array}{c} E \\ p_0(E) > 0 \end{array}}$$

$$\downarrow$$

$$p(\cdot) = p_0(\cdot \mid E)$$

p_0

$$(E_1 : q_1, \ldots, E_k : q_k)$$

$\{E_i\}$ is a partition, $\sum_i q_i = 1$

$$p(\cdot) = \sum_i q_i * p_0(\cdot \mid E_i)$$

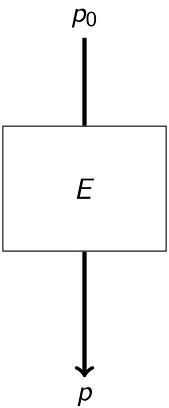
p_0
E
p

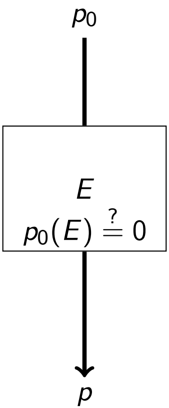
p_0
E
$p_0(E) \stackrel{?}{=} 0$
p

Conditional Probablities

What is the probability that a Democrat will be the next president?

Conditional Probablities

What is the probability that a Democrat will be the next president?

What is the probability that a Democrat will be the next president, *given that a Democrat will be the next president*?

Conditional Probablities

What is the probability that a Democrat will be the next president?

What is the probability that a Democrat will be the next president, *given that a Democrat will be the next president*?

A. Hájek. *What conditional probability could not be*. Synthese, 137, pp. 273 - 323, 2003.

Martingale Property

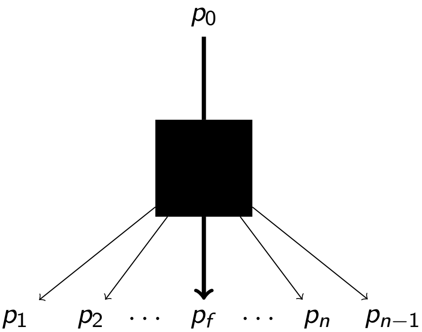

$$(\mathrm{M}) \qquad p_0(A \mid p_f) = p_f(A)$$

"When conditional probability is defined by the ratio rule, it has limited expressive capacity. We would like to allow propositions that have been accorded zero probability to serve as conditions for the probability of other propositions. This is impossible when $p(x \mid a)$ is put as $p(a \wedge x)/p(a)$, for it is undefined when $p(a) = 0$."

D. Makinson. *Conditional Probability in the Light of Qualitative Belief Change*. Journal of Philosophical Logic.

Problem: The condition a is consistent but of zero probability (the **critical zone**).

Problem: The condition a is consistent but of zero probability (the **critical zone**).

Solutions:

- Carnap: If $p(x) = 0$, then x is inconsistent.

Problem: The condition a is consistent but of zero probability (the **critical zone**).

Solutions:

- Carnap: If $p(x) = 0$, then x is inconsistent.

  Define $p_a(\cdot)$ as $p(\cdot \mid a)$. By the left projection, $p_a(x) = p(x \mid a)$, then $p_a(\neg a) = p(\neg a \mid a) = 0$ since $p(a)$. Thus, $p_a(\neg a) = 0$ even though $\neg a$ is inconsistent.

Problem: The condition a is consistent but of zero probability (the **critical zone**).

Solutions:

- Carnap: If $p(x) = 0$, then x is inconsistent.

  Define $p_a(\cdot)$ as $p(\cdot \mid a)$. By the left projection, $p_a(x) = p(x \mid a)$, then $p_a(\neg a) = p(\neg a \mid a) = 0$ since $p(a)$. Thus, $p_a(\neg a) = 0$ even though $\neg a$ is inconsistent.

- $p(x \mid a) = 1$ for every value x when $p(a) = 0$.

Problem: The condition a is consistent but of zero probability (the **critical zone**).

Solutions:

- Carnap: If $p(x) = 0$, then x is inconsistent.

  Define $p_a(\cdot)$ as $p(\cdot \mid a)$. By the left projection, $p_a(x) = p(x \mid a)$, then $p_a(\neg a) = p(\neg a \mid a) = 0$ since $p(a)$. Thus, $p_a(\neg a) = 0$ even though $\neg a$ is inconsistent.

- $p(x \mid a) = 1$ for every value x when $p(a) = 0$. Not very useful.

Problem: The condition a is consistent but of zero probability (the **critical zone**).

Solutions:

- Carnap: If $p(x) = 0$, then x is inconsistent.

  Define $p_a(\cdot)$ as $p(\cdot \mid a)$. By the left projection, $p_a(x) = p(x \mid a)$, then $p_a(\neg a) = p(\neg a \mid a) = 0$ since $p(a)$. Thus, $p_a(\neg a) = 0$ even though $\neg a$ is inconsistent.

- $p(x \mid a) = 1$ for every value x when $p(a) = 0$. Not very useful.

- $p(x \mid a)$ is the limit of the values of $p(x \mid a')$ for suitable infinite sequence of non-critical approximations a' to a. Only defined on special domains.

CPS (Popper Space)

A **conditional probability space** (CPS) over $(W, \mathfrak{A})$ is a tuple $(W, \mathfrak{A}, \mathfrak{B}, \mu)$ such that $\mathfrak{A}$ is an algebra over W, $\mathfrak{B}$ is a set of subsets of W (not necessarily an algebra) that does not contain $\emptyset$ and $\mu : \mathfrak{A} \times \mathfrak{B} \to [0, 1]$ satisfying the following conditions:

1. $\mu(U \mid U) = 1$ if $U \in \mathfrak{B}$
2. $\mu(E_1 \cup E_1 \mid U) = \mu(E_1 \mid U) + \mu(E_2 \mid U)$ if $E_1 \cap E_2 = \emptyset$, $U \in \mathfrak{B}$ and $E_1, E_2 \in \mathfrak{A}$
3. $\mu(E \mid U) = \mu(E \mid X) * \mu(X \mid U)$ if $E \subseteq X \subseteq U$, $U, X \in \mathfrak{B}$ and $E \in \mathfrak{A}$.

$$p : \mathcal{L} \times \mathcal{L} \rightarrow [0, 1]$$

van Fraassen Axioms:

- $vF1$ $p(x, a) = p(x, a')$ whenever $a \equiv a'$
- $vF2$ p_a is a one-place Kolmogorov probability function with $p_a(a) = 1$
- $vF3$ $p(x \wedge y, a) = p(x, a) * p(y, a \wedge x)$ for all a, x, y

"for 'most' values of the right argument of the two-place function, the left projections should be proper one-place Kolmogorov functions, while in the remaining cases it should be the unit function."

(Positive): when $p(a, \top) > 0$ then p_a is a proper Kolmogorov function.

(Carnap) When a is consistent then $p(a, \top) > 0$.

(Unit) When a is consistent but $p(a, \top) = 0$, then p_a is the unit function.

(HL) When a is consistent but $p(a, \top) = 0$, then p_a is a proper Kolmogorov probability function.

What does 'most propositions' mean?

- The van Fraassen system: an unspecified subset (possibly empty) of the consistent propositions,
- The Popper system: all propositions that are above the critical zone or in an unspecified subset (possibly empty) of it,
- The Unit system: for all propositions above the critical zone but no others,
- The Hosiasson-Lindenbaum system: for all propositions above or in the critical zone,
- Carnaps system: we can say any of the last three, since the critical zone is declared empty.

LPS (Lexicographic Probability Space)

A **lexicographic probability space** (LPS) (of length α) is a tuple $(W, \Sigma, \vec{\mu})$ where W is a set of possible worlds, Σ is an algebra over W and $\vec{\mu}$ is a sequence of (finitely/countable additive) probability measures on (W, Σ) indexed by ordinals $< \alpha$.

Fix an *LPS* $\vec{\mu} = (\mu_0, \ldots, \mu_n)$

- E is certain: $\mu_0(E) = 1$

Fix an LPS $\vec{\mu} = (\mu_0, \ldots, \mu_n)$

- E is certain: $\mu_0(E) = 1$
- E is absolutely certain: $\mu_i(E) = 1$ for all $i = 1, \ldots, n$

Fix an *LPS* $\vec{\mu} = (\mu_0, \ldots, \mu_n)$

- E is certain: $\mu_0(E) = 1$
- E is absolutely certain: $\mu_i(E) = 1$ for all $i = 1, \ldots, n$
- E is *assumed*: there exists k such that $\mu_i(E) = 1$ for all $i \leq k$ and $\mu_i(E) = 0$ for all $k < i < n$.

NPS (non-standard probability measures)

$\mathbb{R}^*$ is a *non-Archimedean* field that includes the real numbers as a subfield but also has *infinitesimals*.

For all $b \in \mathbb{R}^*$ such that $-r < b < r$ for some $r \in \mathbb{R}$, there is a unique closest real number a such that $|a - b|$ is an infinitesimal. Let $st(b)$ denote the closest standard real to b.

NPS (non-standard probability measures)

$\mathbb{R}^*$ is a *non-Archimedean* field that includes the real numbers as a subfield but also has *infinitesimals*.

For all $b \in \mathbb{R}^*$ such that $-r < b < r$ for some $r \in \mathbb{R}$, there is a unique closest real number a such that $|a - b|$ is an infinitesimal. Let $st(b)$ denote the closest standard real to b.

A **nonstandard probability space** (NPS) is a tuple (W, Σ, μ) where W is a set of possible worlds, Σ is an algebra over W and μ assigns to elements of Σ, nonnegative elements of $\mathbb{R}^*$ such that $\mu(W) = 1$, $\mu(E \cup F) = \mu(E) + \mu(F)$ if E and F are disjoint.

J. Halpern. *Lexicographic probability, conditional probability, and nonstandard probability*. Games and Economic Behavior, 68:1, pgs. 155 - 179, 2010.

$$
\begin{array}{c}
p_0 \\
\downarrow \\
\boxed{\begin{array}{c} E \\ p_0(E) > 0 \end{array}} \\
\downarrow \\
p(\cdot) = p_0(\cdot \mid E)
\end{array}
$$

p_0

$$(E_1 : q_1, \ldots, E_k : q_k)$$
$\{E_i\}$ is a partition, $\sum_i q_i = 1$

$$p(\cdot) = \sum_i q_i * p_0(\cdot \mid E_i)$$

$$\begin{tikzcd} p_0(\cdot, \top) \arrow[d, dash] \\ \boxed{E} \arrow[d] \\ p(\cdot) = p_0(\cdot, E) \end{tikzcd}$$

p_0

$\mathcal{C}$: set of constraints

p satisfies $\mathcal{C}$

MAXENT

Let us start with the simplest case, where our outcome space, X, contains only a finite number of points, $x_1, x_2, \ldots, x_n$. Then the **entropy** of a probability, p, on this space is:

$$-\sum_i p(x_i) \log p(x_i)$$

and the **information** is the negative of the entropy.

The minimum information or maximum entropy probability is the one which makes the states equiprobable: $p(x_i) = \frac{1}{n}$.

Consider three die x_1, x_2, x_3 and a random variable f such that $f(x_i) = i$.

$$\mathbb{E}[f] = p(x_1)f(x_1) + p(x_2)f(x_2) + p(x_3)f(x_3)$$

What probabilities maximize entropy under the constraint that $\mathbb{E}[f]$ have different values?

MAXENT

$\mathbb{E}[f]$	$p(x_1)$	$p(x_2)$	$p(x_3)$
1	1	0	0
0.1	0.907833	0.084333	0.007834
0.2	0.826297	0.147407	0.026297
$\vdots$	$\vdots$	$\vdots$	$\vdots$
0.8	0.438371	0.323257	0.238271
0.9	0.384586	0.330829	0.284586
2.0	0.333333	0.333333	0.333333
2.1	0.284586	0.330829	0.384586
2.2	0.238372	0.323257	0.438370
$\vdots$	$\vdots$	$\vdots$	$\vdots$
2.8	0.026297	0.147407	0.826296
2.9	0.007834	0.084332	0,907834
3.0	0	0	1

MAXENT

$\mathbb{E}[f]$	$p(x_1)$	$p(x_2)$	$p(x_3)$
1	1	0	0
0.1	0.907833	0.084333	0.007834
0.2	0.826297	0.147407	0.026297
$\vdots$	$\vdots$	$\vdots$	$\vdots$
0.8	0.438371	0.323257	0.238271
0.9	0.384586	0.330829	0.284586
2.0	0.333333	0.333333	0.333333
2.1	0.284586	0.330829	0.384586
2.2	0.238372	0.323257	0.438370
$\vdots$	$\vdots$	$\vdots$	$\vdots$
2.8	0.026297	0.147407	0.826296
2.9	0.007834	0.084332	0,907834
3.0	0	0	1

MAXENT

$\mathbb{E}[f]$	$p(x_1)$	$p(x_2)$	$p(x_3)$
1	1	0	0
0.1	0.907833	0.084333	0.007834
$\vdots$	$\vdots$	$\vdots$	$\vdots$
0.9	0.384586	0.330829	0.284586
2.0	0.333333	0.333333	0.333333
2.1	0.284586	0.330829	0.384586
$\vdots$	$\vdots$	$\vdots$	$\vdots$
2.9	0.007834	0.084332	0.907834
3.0	0	0	1

The MAXENT probabilities are **not** closed under mxing: A mixture of $(1, 0, 0)$ and $(0, 0, 1)$ is $(0.5, 0, 0.5)$, but this is not in the list...

Kullback-Leibler

Suppose that we start with a prior probability, p_0, and move to a posterior p_1 which satisfies certain constraints. The Kullback-Leibler "distance" is:

$$I(p_1, p_0) = \sum_i p_1(x_i) \log \frac{p_1(x_i)}{p_0(x_i)}$$

p_0

$\mathcal{C}$: set of constraints

p satisfies $\mathcal{C}$

p minimizes $I(p, p_0)$

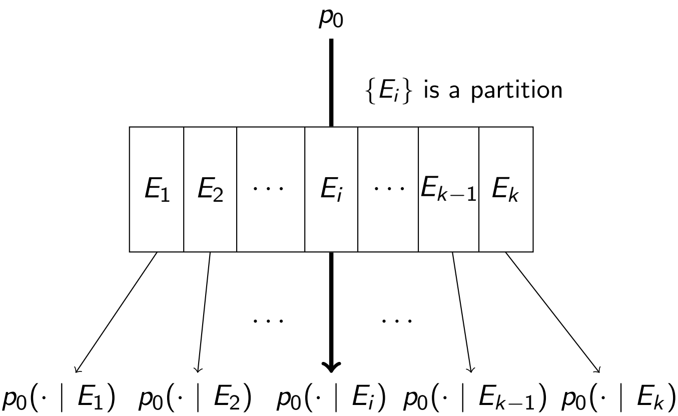
p_0
$\{E_i\}$ is a partition
E_1 E_2 $\cdots$ E_i $\cdots$ E_{k-1} E_k
$\cdots$ $\cdots$
$p_0(\cdot \mid E_1)$ $p_0(\cdot \mid E_2)$ $p_0(\cdot \mid E_i)$ $p_0(\cdot \mid E_{k-1})$ $p_0(\cdot \mid E_k)$

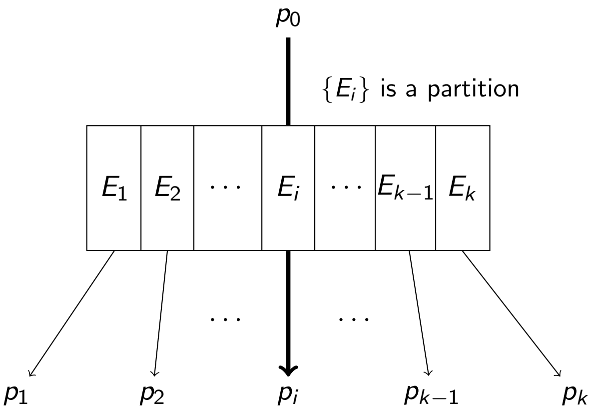
p_0
$\{E_i\}$ is a partition
E_1 E_2 $\cdots$ E_i $\cdots$ E_{k-1} E_k
$\cdots$ $\cdots$
p_1 p_2 p_i p_{k-1} p_k

Suppose that you are in a learning situation even more amorphous than the kind which motivates Jeffrey's idea. There is no nontrivial partition that you expect with probability one to be sufficient for your belief change....Perhaps you are in a novel situation where you expect the unexpected observational input....You are going to just think about some subject matter and update as a result of your thoughts...I will consider the learning situation a kind of black box and attempt no analysis of its internal structure.

(Skyrms, pg. 96, 97)

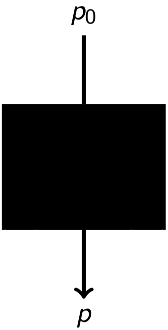
p_0
p

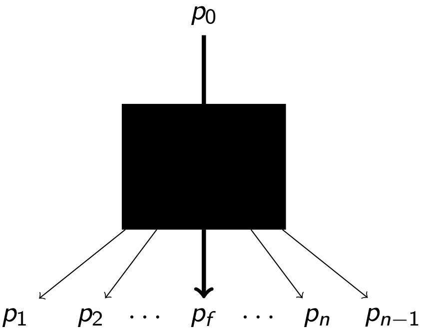
p_0
p_1
p_2 $\cdots$ p_f $\cdots$ p_n p_{n-1}

It was suggested by Skyrms (1990) that this principle provides a plausible way to distinguish learning situations from situations where one expects probabilities to change for other reasons, such as getting drunk, having a brain lesion or having a dangerously low blood sugar level.

It was suggested by Skyrms (1990) that this principle provides a plausible way to distinguish learning situations from situations where one expects probabilities to change for other reasons, such as getting drunk, having a brain lesion or having a dangerously low blood sugar level.

Huttegger develops an account in which the reflection principle is a necessary condition for a black-box probability update to count as a *genuine learning experience*.

Simon Huttegger. *Learning Experiences and the Value of Knowledge*. Philosophical Studies, 2013.

Plan

- Day 1 Introduction to belief revision, AGM, possible worlds models, Bayesian models (time permitted)
- Day 2 Bayesian models (continued), Justifying Bayesian models (Dutch books, Accuracy-based arguments), Updating probabilities
- Day 3 The value of learning, Lottery Paradox, Preface Paradox, Review Paradox, Iterated belief revision, Context shifts, Becoming aware
- Day 4 The value of learning, Lottery Paradox, Preface Paradox, Review Paradox, Iterated belief revision, Context shifts, Becoming aware (continued)
- Day 5 Interactive epistemology (Agreement Theorems, Belief Revision in Games)

Topics

- Classic papers (Makinson, Diaconis & Zabel, KLM, ...)
- Beliefs, credences and probability (Leitgeb's stability theory of belief, Pettigrew, Fitelson & Shear)
- Revising probabilities (List, Dietrich & Bradley, Halpern)
- Conditioning vs. learning (Osherson et al., Curpi et al.)
- Context shifts (Halpern & Grünwald, Romeijn, Pettigrew)
- Lottery, Preface and Review paradox (Leitgeb, Easerwen & Fitelson)
- Iterated belief change, long-term dynamics, convergence results (Huttegger, EP)
- Bayesian reasoning, reasoning to the best explanation, case-base reasoning (Gilboa et al., Douven and Shubach)