Logic and Probabilistic Models of Belief Change

Eric Pacuit

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

February 20, 2016

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

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

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

$$p_0 \quad \Longrightarrow \quad p_t$$

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

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

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

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 $\models \neg(\varphi \wedge \psi)$.

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

H. Greaves. *Epistemic decision theory*. Mind (2013.

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

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

"The ratio is an analysis not a definition of conditional probability".

Bayes Theorem

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

"The ratio is an analysis not a definition of conditional probability".

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

Bayes Theorem

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

"The ratio is an analysis not a definition of conditional probability".

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*?

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

$$q(H) = \sum_i p(H \mid E_i) q(E_i)$$

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

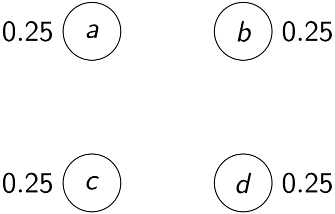

$$p(E_1) = 0.8$$

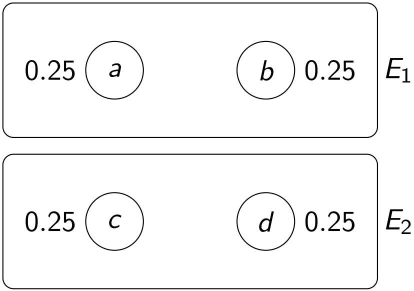

The probability that the guilty party is left-handed is 0.8

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

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

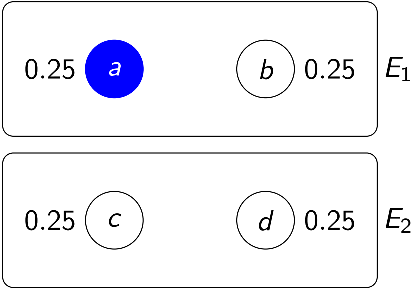

The probability that the guilty party is left-handed is 0.8

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

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

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

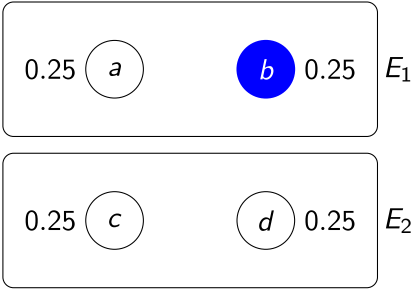

The probability that the guilty party is left-handed is 0.8

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

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

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

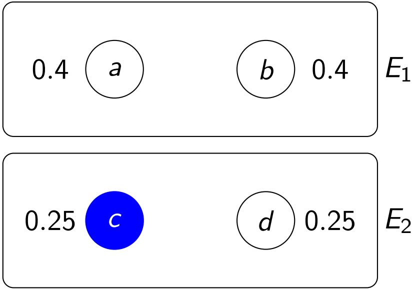

The probability that the guilty party is left-handed is 0.8

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

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

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

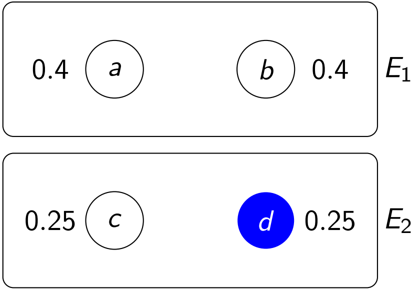

The probability that the guilty party is left-handed is 0.8

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

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

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

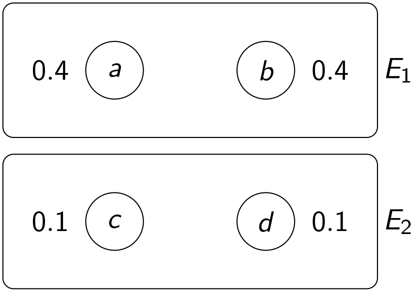

The probability that the guilty party is left-handed is 0.8

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

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

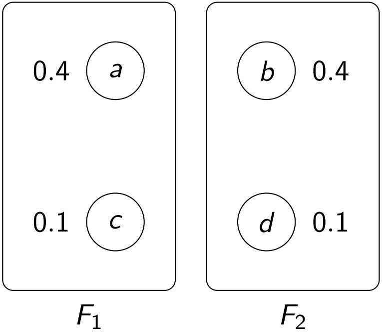

The probability that the guilty party is left-handed is 0.8

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

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

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

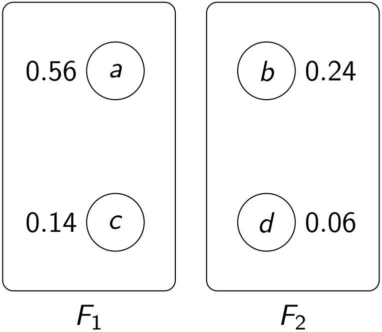

The probability that the guilty party is left-handed is 0.8

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

$$p(F_1) = 0.7 \qquad p(E_2) = 0.3$$

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.

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

Fact. Jeffrey conditioning is not commutative.

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

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.

"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 $p(a) = 0$.

- Borel: Suppose a point is selected at random from the surface of the earth. What is the probability that it lies in the Western hemisphere, given that it lies on the equator?

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: Whenever $p(x) = 0$ then x is inconsistent.

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

Solutions:

- Carnap: Whenever $p(x) = 0$ 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: Whenever $p(x) = 0$ 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: Whenever $p(x) = 0$ 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: Whenever $p(x) = 0$ 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 $E \in \mathfrak{A}$.

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

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.
  m

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.

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

Bridge Principles

Probability 1: $Bel(A)$ iff $P(A) = 1$

Bridge Principles

Probability 1: $Bel(A)$ iff $P(A) = 1$

The Lockean Thesis: $Bel(A)$ iff $P(A) > r$

Bridge Principles

Probability 1: $Bel(A)$ iff $P(A) = 1$

The Lockean Thesis: $Bel(A)$ iff $P(A) > r$

Decision-theoretic accounts: $Bel(A)$ iff
$\sum_{w \in W} P(\{w\}) \cdot u(bel\ A, w)$ has such-and-such property

Bridge Principles

Probability 1: $Bel(A)$ iff $P(A) = 1$

The Lockean Thesis: $Bel(A)$ iff $P(A) > r$

Decision-theoretic accounts: $Bel(A)$ iff
$\sum_{w \in W} P(\{w\}) \cdot u(bel\ A, w)$ has such-and-such property

The Nihilistic proposal: "...no explication of belief is possible within the confines of the probability model."

Preface Paradox

D. Makinson. *The Paradox of the Preface*. Analysis, 25, 205 - 207, 1965.

Preface Paradox

Suppose that in the course of his book an author makes a great many assertions: $s_1, s_2, \ldots, s_n$.

Preface Paradox

Suppose that in the course of his book an author makes a great many assertions: $s_1, s_2, \ldots, s_n$.

Given each one of these, he believes that it is true (for each i, $B_A(s_i)$)

Preface Paradox

Suppose that in the course of his book an author makes a great many assertions: $s_1, s_2, \ldots, s_n$.

Given each one of these, he believes that it is true (for each i, $B_A(s_i)$)

If he has already written other books, and received corrections from readers and reviewers, he may also believe that not everything he has written in his latest book is true.

Preface Paradox

Suppose that in the course of his book an author makes a great many assertions: $s_1, s_2, \ldots, s_n$.

Given each one of these, he believes that it is true (for each i, $B_A(s_i)$)

If he has already written other books, and received corrections from readers and reviewers, he may also believe that not everything he has written in his latest book is true.

$B_A(\neg(s_1 \wedge s_2 \wedge \cdots \wedge s_n))$

Preface Paradox

Suppose that in the course of his book an author makes a great many assertions: $s_1, s_2, \ldots, s_n$.

Given each one of these, he believes that it is true (for each i, $B_A(s_i)$)

If he has already written other books, and received corrections from readers and reviewers, he may also believe that not everything he has written in his latest book is true.

$B_A(\neg(s_1 \wedge s_2 \wedge \cdots \wedge s_n))$

But $\{s_1, \ldots, s_n, \neg(s_1 \wedge \cdots \wedge s_n)\}$ is logically inconsistent.

Preface Paradox

A philosopher who asserts "all of my present philosophical positions are correct" would be regarded as rash and over-confident

A philosopher who asserts "at least some of my present philosophical beliefs will turn out to be incorrect" is simply being sensible and honest.

Preface Paradox

1. each belief from the set $\{s_1, \ldots, s_n, s_{n+1}\}$ is rational
2. the set $\{s_1, \ldots, s_n, s_{n+1}\}$ of beliefs is rational.

1. does not necessarily imply 2.

Preface Paradox: The Problem

"The author of the book is being rational even though inconsistent. More than this: he is being rational even though he believes each of a certain collection of statements, which *he knows* are logically incompatible....this appears to present a living and everyday example of a situation which philosophers have commonly dismissed as absurd; that it is sometimes rational to hold incompatible beliefs."

D. Makinson. *The Paradox of the Preface*. Analysis, 25, 205 - 207, 1965.

H. Leitgeb. *The Review Paradox: On the Diachronic Costs of Not Closing Rational Belief Under Conjunction*. Nous, 2013.

Bel_t is the set of propositions believed at time t

P_t is the agent's degree of belief function at time t

$t' > t$

P1 If the degrees of belief that the agents assigns to two propositions are identical then either the agent believes both of them or neither of them.

For all X, Y: if $P_t(X) = P_t(Y)$, then $Bel_t(X)$ iff $Bel_t(Y)$.

P2 If the agent already believes X, then updating on the piece of evidence X does not change her system of (all-or-nothing) beliefs at all.

For all X: if the evidence that the agent obtains between t and $t' > t$ is the proposition X, but it holds already that $Bel_t(X)$, then for all Y:

$$Bel_{t'}(Y) \text{ iff } Bel_t(Y)$$

P3 When the agent learns, this is captured probabilistically by conditionalization.

For all X (with $P_t(X) > 0$): if the evidence that the agent obtains between t and $t' > t$ is the proposition X, but it holds already that $Bel_t(X)$, then for all Y:

$$P_{t'}(Y) = P_t(Y \mid X)$$

Assume $Bel_t(A), Bel_t(B)$ but not $Bel_t(A \cap B)$

- Suppose that the agent receive A as evidence.
- $P_{t'}(B) = P_t(B \mid A) = P_t(A \cap B \mid A) = P_{t'}(A \cap B)$.
- By P1, the agent must have the same doxastic attitude towards B and $A \cap B$.
- By P2, the agent's attitude towards B and $A \cap B$ must be the same at t' as at t.
- But, $Bel_t(B)$ and not $Bel_t(A \cap B)$

$Bel_t(A), Bel_t(B)$

$\neg Bel_t(A \cap B)$

$0 < P_t(A) < 1$

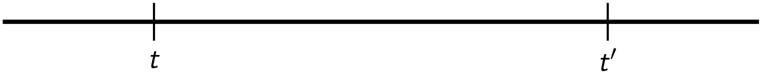

Assumption

$Bel_t(A), Bel_t(B)$

$\neg Bel_t(A \cap B)$

$0 < P_t(A) < 1$

t Receives evidence A t'

$P_{t'}(B) = P_t(B \mid A)$

$P_{t'}(A \cap B) = P_t(A \cap B \mid A) = P_t(B \mid A)$

By P3

$Bel_t(A), Bel_t(B)$

$\neg Bel_t(A \cap B)$

$0 < P_t(A) < 1$

$P_{t'}(B) = P_t(B \mid A)$

$P_{t'}(A \cap B) = P_t(A \cap B \mid A) = P_t(B \mid A)$

$Bel_{t'}(B)$ iff $Bel_{t'}(A \cap B)$

t — Receives evidence A — t'

By P1

$Bel_t(A)$, $Bel_t(B)$

$\neg Bel_t(A \cap B)$

$0 < P_t(A) < 1$

$P_{t'}(B) = P_t(B \mid A)$

$P_{t'}(A \cap B) = P_t(A \cap B \mid A) = P_t(B \mid A)$

$Bel_{t'}(B)$ iff $Bel_{t'}(A \cap B)$

t — Receives evidence A — t'

$Bel_t(A)$ iff $Bel_{t'}(A)$

$Bel_t(B)$ iff $Bel_{t'}(B)$

$Bel_t(A \cap B)$ iff $Bel_{t'}(A \cap B)$

By P2

$Bel_t(B)$

$\neg Bel_t(A \cap B)$

$0 < P_t(A) < 1$

$P_{t'}(B) = P_t(B \mid A)$

$P_{t'}(A \cap B) = P_t(A \cap B \mid A) = P_t(B \mid A)$

$Bel_{t'}(B)$ iff $Bel_{t'}(A \cap B)$

t Receives evidence A t'

$Bel_t(A)$ iff $Bel_{t'}(A)$

$Bel_t(B)$ iff $Bel_{t'}(B)$

$Bel_t(A \cap B)$ iff $Bel_{t'}(A \cap B)$

$Bel_t(B)$ iff $Bel_{t'}(B)$ iff $Bel_{t'}(A \cap B)$ iff $Bel_t(A \cap B)$

Lottery Paradox

H. Kyburg. *Probability and the Logic of Rational Belief*. Wesleyan University Press, 1961.

I. Douven and T. Williamson. *Generalizing the Lottery Paradox*. British Journal of the Philosophy of Science, 57, 755 - 779, 2006.

G. Wheeler. *A Review of the Lottery Paradox*. Probability and Inference: Essays in honor of Henry E. Kyburg, Jr., College Publications, 2007.

Lottery Paradox

Lottery Paradox

Consider a fair lottery with 1,000,000 tickets and one prize.

Lottery Paradox

Consider a fair lottery with 1,000,000 tickets and one prize.

The probability that a given ticket will win is 0.000001 $(1/1,000,000)$ and the probability that it will not win is 0.999999.

Lottery Paradox

Consider a fair lottery with 1,000,000 tickets and one prize.

The probability that a given ticket will win is 0.000001 $(1/1,000,000)$ and the probability that it will not win is 0.999999.

"Surely if a sheer probability is ever sufficient to warrant the acceptance of a hypothesis, this is a case"

Lottery Paradox

Consider a fair lottery with 1,000,000 tickets and one prize.

The probability that a given ticket will win is 0.000001 $(1/1,000,000)$ and the probability that it will not win is 0.999999.

"Surely if a sheer probability is ever sufficient to warrant the acceptance of a hypothesis, this is a case"

For each lottery ticket t_i ($i = 1, \ldots, 1000000$), the agent believes that t_i will loose $B_A(\neg$'t_i will win')

Lottery Paradox

A **rule of acceptance**: If S and T are acceptable statements, their conjunction is also acceptable.

Lottery Paradox

A **rule of acceptance**: If S and T are acceptable statements, their conjunction is also acceptable.

So, the conjunction $\bigwedge_{i=1}^{1000000}$ 't_i will not win' should be accepted. That is, the agent should rationally accept that *no lottery ticket will win*.

Lottery Paradox

A **rule of acceptance**: If S and T are acceptable statements, their conjunction is also acceptable.

So, the conjunction $\bigwedge_{i=1}^{1000000}$ 't_i will not win' should be accepted. That is, the agent should rationally accept that *no lottery ticket will win*.

But, this is a fair lottery, so at least one ticket is *guaranteed* to win!

The Lottery Paradox

Kyburg: The following are inconsistent,

1. It is rational to accept a proposition that is very likely true,
2. It is not rational to accept a propositional that you are aware is inconsistent
3. It is rational to accept a proposition P and it is rational to accept another proposition P' then it is rational to accept $P \wedge P'$