Logic and Probabilistic Models of Belief Change

Eric Pacuit

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

February 12, 2016

$$K \circ \varphi$$

$$K \circ \varphi$$

Initial set of beliefs

$K \circ \varphi$

Initial set of beliefs

New evidence φ

Belief change operator: $\circ : \mathcal{B} \times \mathcal{L} \to \mathcal{B}$

$$K \circ \varphi$$

Initial set of beliefs

New evidence φ

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

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

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

Bayes theorem is important because it expresses the quantity $p(E|F)$ (the probability of a hypothesis E given the evidence F) —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.

Bayes theorem is important because it expresses the quantity $p(E|F)$ (the probability of a hypothesis E given the evidence F) —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. That is, compare $p(Dice \mid Twelve)$ with $p(Roulette \mid Twelve)$.

Bayes theorem is important because it expresses the quantity $p(E|F)$ (the probability of a hypothesis E given the evidence F) —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. 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 $p(\mathit{Twelve} \mid \mathit{Roulette}) = 1/38$.

Bayes theorem is important because it expresses the quantity $p(E|F)$ (the probability of a hypothesis E given the evidence F) —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. 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 $p(\mathit{Roulette})$.

Bayes theorem is important because it expresses the quantity $p(E|F)$ (the probability of a hypothesis E given the evidence F) —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. 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 $p(Twelve \mid Roulette) = 1/38$. Based on our observations about the casino, we can judge the prior probabilities $p(Dice)$ and $p(Roulette)$. 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: $Bel : A \rightarrow [0, 1]$ 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: $\mu : \Sigma \rightarrow \mathbb{R}^*$

- Dempster-Shafer belief functions: $Bel : A \rightarrow [0,1]$ 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: $\mu : \Sigma \rightarrow \mathbb{R}^*$

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

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**: $\sigma^*(A) = \sup\{x \mid x \in \sigma(A)\}$

Signals/Knowledge/Questions/etc.

Epistemic Probability Models

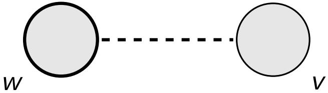

$\mathcal{M} = \langle W, \{\Pi_i\}_{i \in \mathcal{A}} \rangle$
Π_i is agent i's partition with $\Pi_i(w)$ the partition cell containing w.

$K_i(E) = \{w \mid \Pi_i(w) \subseteq E\}$

Epistemic Probability Models

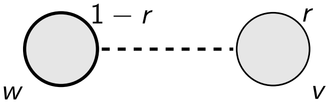

$\mathcal{M} = \langle W, \{\Pi_i\}_{i\in\mathcal{A}}, \{p_i\}_{i\in\mathcal{A}}\rangle$
for each i, $p_i : W \to [0,1]$ is a probability measure

$$B_i^r(E) = \{w \mid p_i(E \mid \Pi_i(w)) = \frac{p_i(E \cap \Pi_i(w))}{p_i(\Pi_i(w))} \geq r\}$$

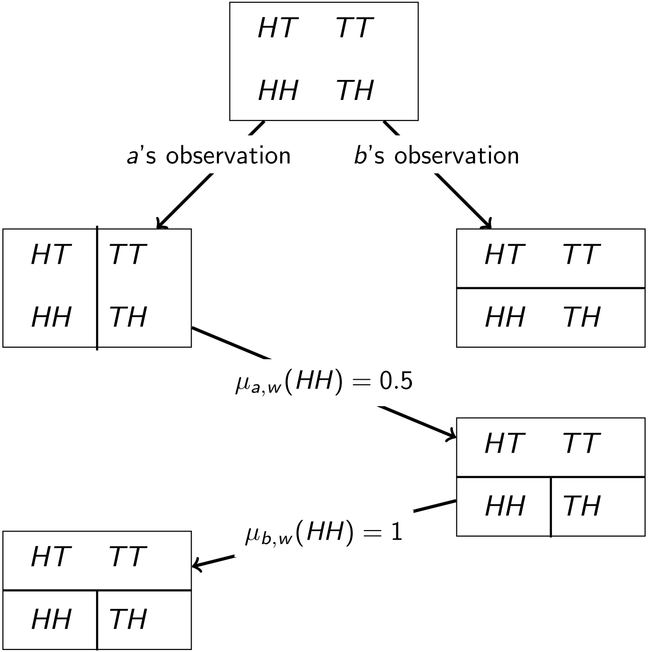
HT TT
HH TH
a's observation
b's observation
HT TT
HH TH
HT TT
HH TH
$\mu_{a,w}(HH) = 0.5$
HT TT
HH TH
$\mu_{b,w}(HH) = 1$
HT TT
HH TH

2 Scientists Perform an Experiment

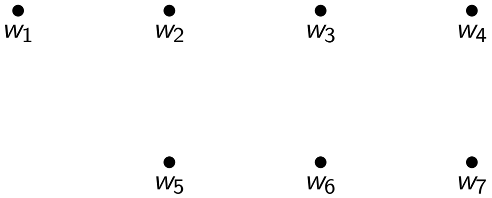

They agree the true state is one of seven different states.

2 Scientists Perform an Experiment

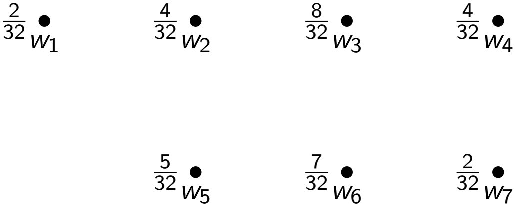

They agree on a common prior.

2 Scientists Perform an Experiment

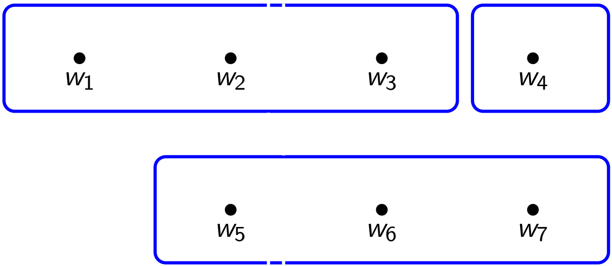

They agree that Experiment 1 would produce the blue partition.

2 Scientists Perform an Experiment

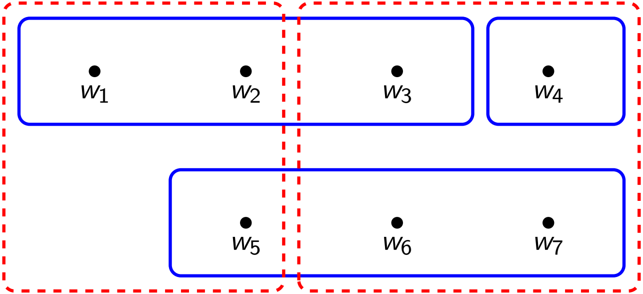

They agree that Experiment 1 would produce the blue partition and Experiment 2 the red partition.

2 Scientists Perform an Experiment

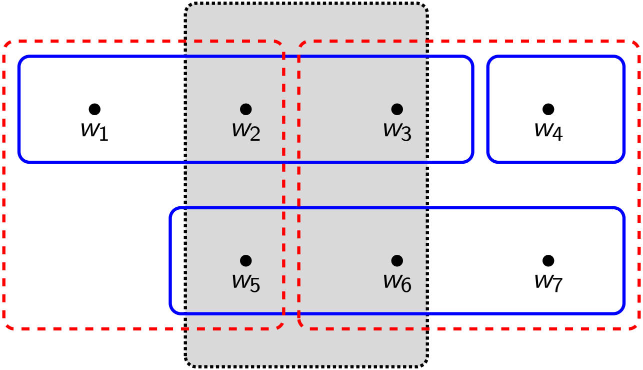

They are interested in the truth of $E = \{w_2, w_3, w_5, w_6\}$.

2 Scientists Perform an Experiment

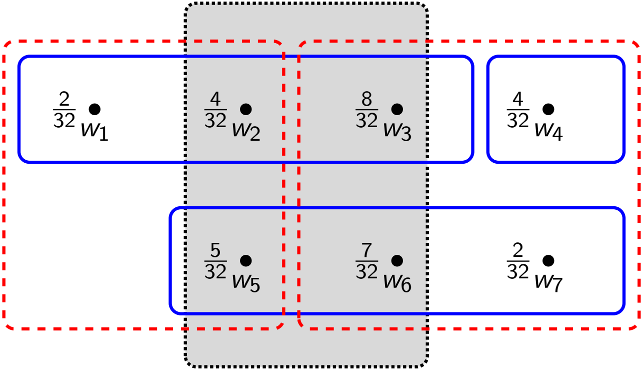

So, they agree that $P(E) = \frac{24}{32}$.

2 Scientists Perform an Experiment

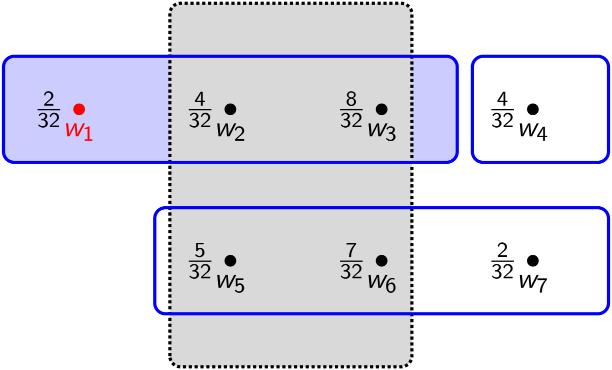

Also, that *if the true state is* w_1, then Experiment 1 will yield

$$P(E|I) = \frac{P(E \cap I)}{P(I)} = \frac{12}{14}$$

2 Scientists Perform an Experiment

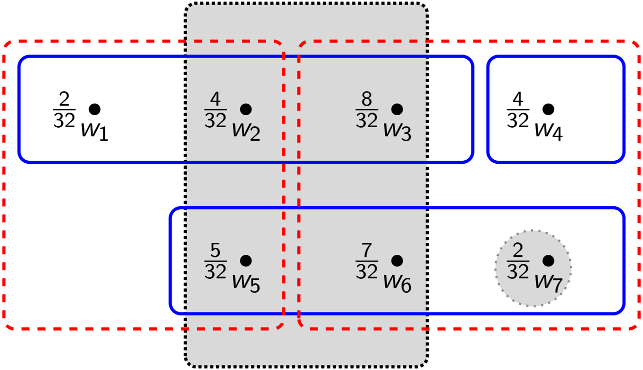

Suppose the true state is w_7 and the agents preform the experiments.

2 Scientists Perform an Experiment

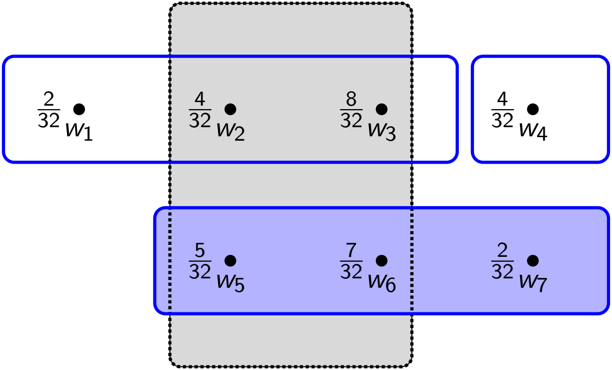

Suppose the true state is w_7, then $Pr_1(E) = \frac{12}{14}$

2 Scientists Perform an Experiment

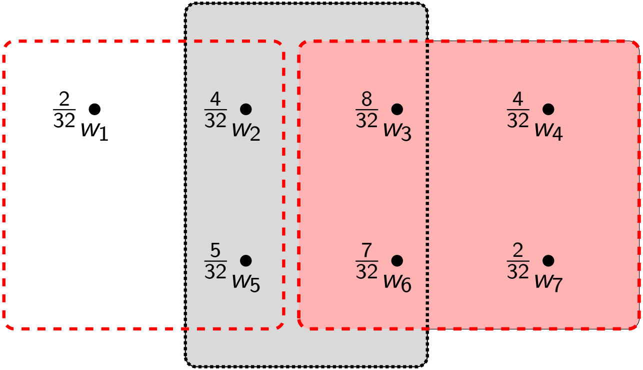

Then $Pr_1(E) = \frac{12}{14}$ and $Pr_2(E) = \frac{15}{21}$

2 Scientists Perform an Experiment

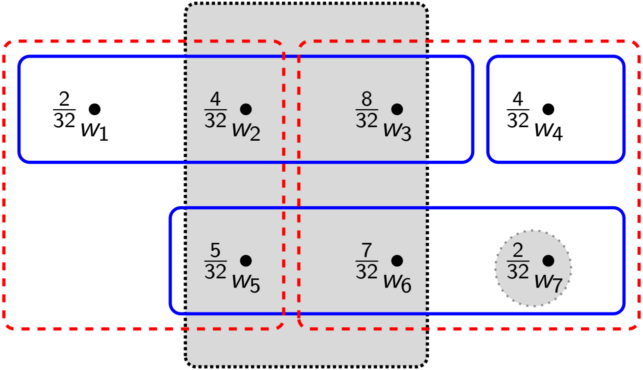

Suppose they exchange emails with the new subjective probabilities: $Pr_1(E) = \frac{12}{14}$ and $Pr_2(E) = \frac{15}{21}$

2 Scientists Perform an Experiment

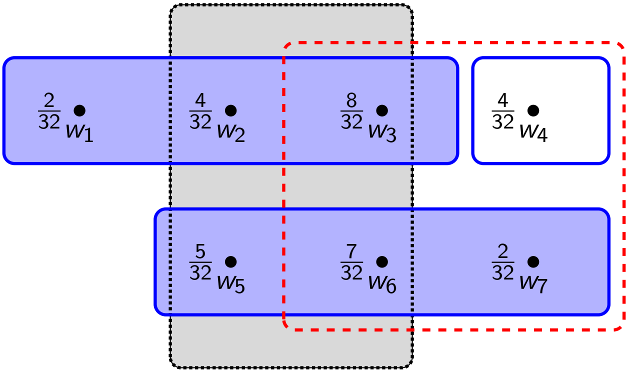

Agent 2 learns that w_4 is **NOT** the true state (same for Agent 1).

2 Scientists Perform an Experiment

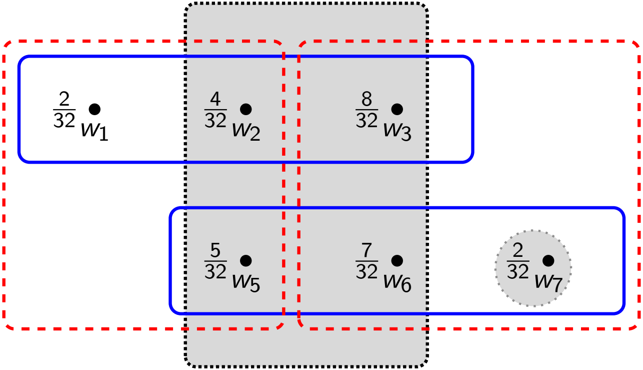

Agent 2 learns that w_4 is **NOT** the true state (same for Agent 1).

2 Scientists Perform an Experiment

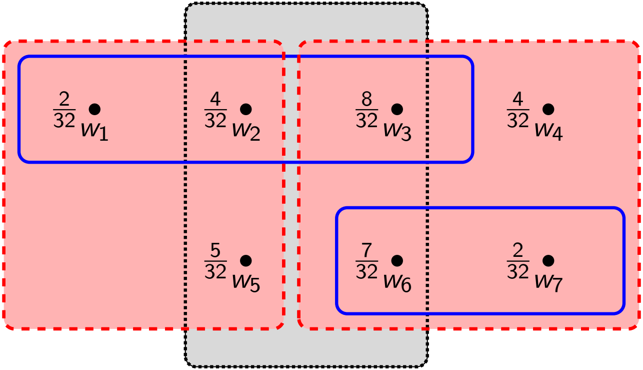

Agent 1 learns that w_5 is **NOT** the true state (same for Agent 1).

2 Scientists Perform an Experiment

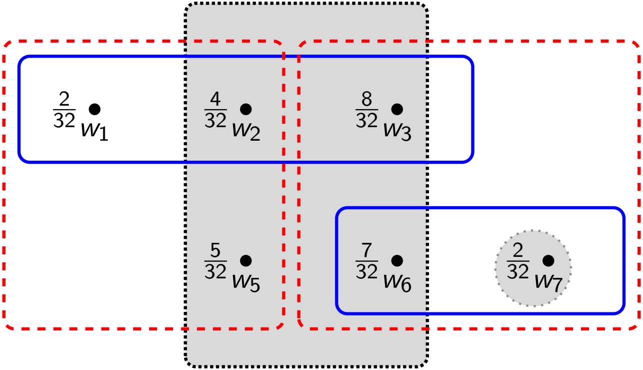

The new probabilities are $Pr_1(E|I') = \frac{7}{9}$ and $Pr_2(E|I') = \frac{15}{17}$

2 Scientists Perform an Experiment

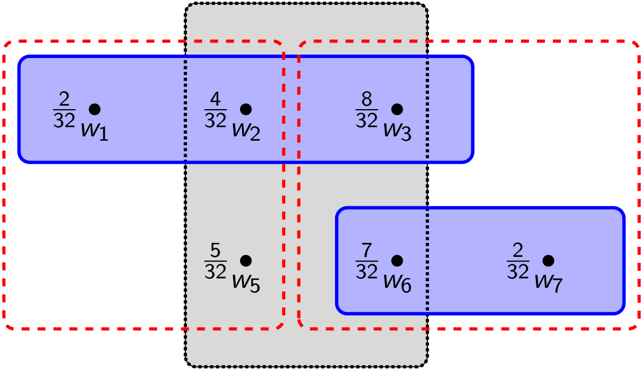

After exchanging this information ($Pr_1(E|I') = \frac{7}{9}$ and $Pr_2(E|I') = \frac{15}{17}$), Agent 2 learns that w_3 is **NOT** the true state.

2 Scientists Perform an Experiment

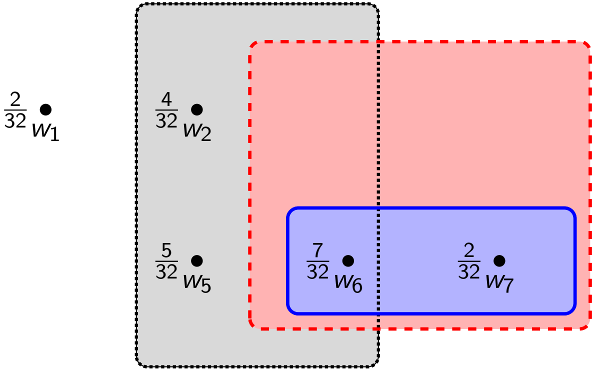

No more revisions are possible and the agents agree on the posterior probabilities.

1. Belief, credence and probability: Dutch book, Lottery Paradox, Preface Paradox, Review Paradox
2. Learning vs. supposing/Naive vs. sophisticated spaces
3. Justifying conditionalization

Dutch Book Arguments

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

J. Joyce. *Bayesianism*. in Handbook of Rationality.

Ann: "the probability it will rain tomorrow is 0.9" means "Ann's degree of belief is fairly high (0.9) that it will rain tomorrow. Of course whether it will actually rain, depends on objective events taking place in the external worlds."

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 if X, b 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 ($X \cap Y = \emptyset$)

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

Dutch Book

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

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

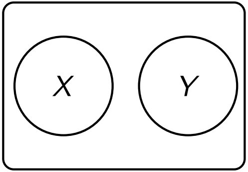

Dutch Book

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

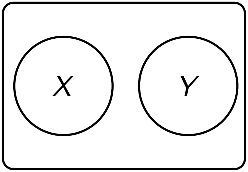

Dutch Book

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 $\mathcal{W}_2 = \{0.5, W_{X \vee Y}\}$

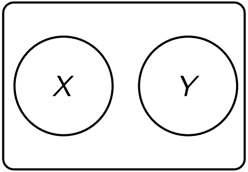

Dutch Book

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 $\mathcal{W}_2$

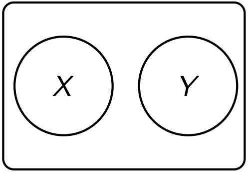

Dutch Book

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 $\mathcal{W}_2$

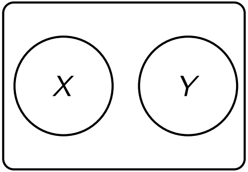

Dutch Book

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 $\mathcal{W}_2 = \{0.5, W_{X\vee Y}\}$

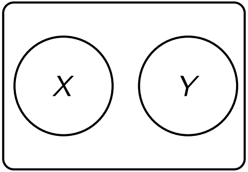

- indifferent between $\mathcal{W}_1$ and $\mathcal{W}_2$
- swap $\mathcal{W}_1$ for $\mathcal{W}_2$
- But $\mathcal{W}_1$ is always better:

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 $\mathcal{W}_2 = \{0.5, W_{X \vee Y}\}$

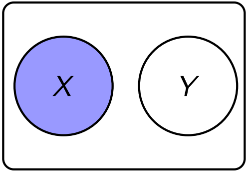

- indifferent between $\mathcal{W}_1$ and $\mathcal{W}_2$
- swap $\mathcal{W}_1$ for $\mathcal{W}_2$
- But $\mathcal{W}_1$ is always better:
  - If X is true
    payoff($\mathcal{W}_1$) $= 1.6 >$ payoff($\mathcal{W}_2$)$=1.5$

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 $\mathcal{W}_2 = \{0.5, W_{X \vee Y}\}$

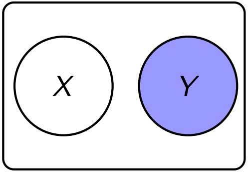

- indifferent between $\mathcal{W}_1$ and $\mathcal{W}_2$
- swap $\mathcal{W}_1$ for $\mathcal{W}_2$
- But $\mathcal{W}_1$ is always better:
  - If X is true
    payoff($\mathcal{W}_1$) $= 1.6 >$ payoff($\mathcal{W}_2$)$=1.5$
  - If Y is true
    payoff($\mathcal{W}_1$) $= 1.6 >$ payoff($\mathcal{W}_2$)$=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

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 $\mathcal{W}_2 = \{0.5, W_{X \vee Y}\}$

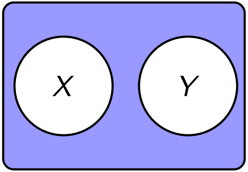

- indifferent between $\mathcal{W}_1$ and $\mathcal{W}_2$
- swap $\mathcal{W}_1$ for $\mathcal{W}_2$
- But $\mathcal{W}_1$ is always better:
    - If X is true
      payoff($\mathcal{W}_1$) $= 1.6 >$ payoff($\mathcal{W}_2$)$=1.5$
    - If Y is true
      payoff($\mathcal{W}_1$) $= 1.6 >$ payoff($\mathcal{W}_2$)$=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

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

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