Decision and Social Choice

PE3101P — Lecture 2: Probability

Tomi Francis

National University of Singapore

20th August, 2026

Today

The week 1 exercise sheet is up

- On my website: `https://tomifrancis.com/pe3101p.html`

The week 1 exercise sheet is up

- On my website: `https://tomifrancis.com/pe3101p.html`
- Optional, and not for credit.

A new document

I'm giving you the greatest gift a student can receive:

A new document

I'm giving you the greatest gift a student can receive:

extra homework.

A new document

I'm giving you the greatest gift a student can receive:

extra homework.

- A **set theory primer**, on the website.
- Optional — but may be useful for the for-credit assignment and other future material

A new document

I'm giving you the greatest gift a student can receive:

extra homework.

- A **set theory primer**, on the website.
- Optional — but may be useful for the for-credit assignment and other future material
- In pdf and html. **I recommend the html**.

First for-credit assignment

- Goes out on my website **over the weekend**.

First for-credit assignment

- Goes out on my website **over the weekend**.
- Due **Friday of Week 4 — 4th September**, by midnight.

My office

- I'm moving office today. From now on I'm in **AS3-05-01**.

My office

- I'm moving office today. From now on I'm in **AS3-05-01**.
- **Office hours today are 16:30–18:00**, later than usual — I'm moving in right after this lecture.

The register

- Printed this time. **Find your name and sign.**

The register

- Printed this time. **Find your name and sign.**
- It's in alphabetical order.

The register

- Printed this time. **Find your name and sign.**
- It's in alphabetical order.
- Can't find yourself? Then I don't have records of you yet — add your name and student number, then sign.

Two senses of “ought”

Fact-relative

What will *in fact* be best.

Two senses of "ought"

Fact-relative

What will *in fact* be best.

Evidence-relative

What your **evidence indicates** will be best, on the balance of probabilities.

Two senses of "ought"

Fact-relative

What will *in fact* be best.

Evidence-relative

What your **evidence indicates** will be best, on the balance of probabilities.

- Run across the road without looking and survive: correct in the fact-relative sense, a mistake in the evidence-relative one.

Two senses of "ought"

Fact-relative

What will *in fact* be best.

Evidence-relative

What your **evidence indicates** will be best, on the balance of probabilities.

- Run across the road without looking and survive: correct in the fact-relative sense, a mistake in the evidence-relative one.
- **Decision theory is about the evidence-relative sense.**

Acts, states, outcomes

	Miners in A	Miners in B
Block A	all 10 live	all 10 die
Block B	all 10 die	all 10 live
Block neither	1 dies	1 dies

Acts, states, outcomes

	Miners in A	Miners in B
Block A	all 10 live	all 10 die
Block B	all 10 die	all 10 live
Block neither	1 dies	1 dies

- **Acts**: the rows — the decisions you could make.

Acts, states, outcomes

	Miners in A	Miners in B
Block A	all 10 live	all 10 die
Block B	all 10 die	all 10 live
Block neither	1 dies	1 dies

- **Acts**: the rows — the decisions you could make.
- **States of nature**: the columns — ways the world could be, which you're uncertain about.

Acts, states, outcomes

	Miners in A	Miners in B
Block A	all 10 live	all 10 die
Block B	all 10 die	all 10 live
Block neither	1 dies	1 dies

- **Acts**: the rows — the decisions you could make.
- **States of nature**: the columns — ways the world could be, which you're uncertain about.
- **Outcomes**: the cells — how things turn out.

Conditions on the states — what I gave you

(i) Together with the act chosen, they **resolve the uncertainty** and guarantee one particular outcome.

Conditions on the states — what I gave you

(i) Together with the act chosen, they **resolve the uncertainty** and guarantee one particular outcome.

(ii) The act you choose **does not affect** whether the state occurs, or how probable it is.

Two more conditions

These two are about the **set** of states, not each state on its own.

Two more conditions

These two are about the **set** of states, not each state on its own.

(iii) **Mutually exclusive** — at most one can happen.

Two more conditions

These two are about the **set** of states, not each state on its own.

(iii) **Mutually exclusive** — at most one can happen.

(iv) **Jointly exhaustive** — at least one must happen.

Two more conditions

These two are about the **set** of states, not each state on its own.

(iii) **Mutually exclusive** — at most one can happen.
(iv) **Jointly exhaustive** — at least one must happen.

So **exactly one** state is always the actual one.

Two more conditions

These two are about the **set** of states, not each state on its own.

(iii) **Mutually exclusive** — at most one can happen.

(iv) **Jointly exhaustive** — at least one must happen.

So **exactly one** state is always the actual one.

Like rolling a die: one number ends up face up.

Chance and credence

- Which act is sensible depends on **how likely** you think each state is. That's your **credence**.

Chance and credence

- Which act is sensible depends on **how likely** you think each state is. That's your **credence**.
- Credences are a lot like chances, and we call both probabilities.

Chance and credence

- Which act is sensible depends on **how likely** you think each state is. That's your **credence**.
- Credences are a lot like chances, and we call both probabilities.
- But **objective chances** are different: they're out in the world.

Chance and credence

- Which act is sensible depends on **how likely** you think each state is. That's your **credence**.
- Credences are a lot like chances, and we call both probabilities.
- But **objective chances** are different: they're out in the world.
- Frequencies, or what a theory like quantum mechanics puts out.

Chance and credence

- Which act is sensible depends on **how likely** you think each state is. That's your **credence**.
- Credences are a lot like chances, and we call both probabilities.
- But **objective chances** are different: they're out in the world.
- Frequencies, or what a theory like quantum mechanics puts out.
- **Credences are in the agent's head** — how likely she thinks something is, given her evidence.

Utility and expected utility

A fair coin. **Take the bet**: \$100 on heads, nothing on tails.
Decline: \$50 either way.

	Heads ($p = 0.5$)	Tails ($p = 0.5$)
Take the bet	100	0
Decline	50	50

Utility and expected utility

A fair coin. **Take the bet**: \$100 on heads, nothing on tails. **Decline**: \$50 either way.

	Heads ($p = 0.5$)	Tails ($p = 0.5$)
Take the bet	100	0
Decline	50	50

- **Utilities** are numbers for how good an outcome is, by the agent's lights.

Utility and expected utility

A fair coin. **Take the bet**: \$100 on heads, nothing on tails. **Decline**: \$50 either way.

	Heads ($p = 0.5$)	Tails ($p = 0.5$)
Take the bet	100	0
Decline	50	50

- **Utilities** are numbers for how good an outcome is, by the agent's lights.
- Where do the numbers come from? Later in the course. For now: **more is better**.

Utility and expected utility

A fair coin. **Take the bet**: \$100 on heads, nothing on tails. **Decline**: \$50 either way.

	Heads ($p = 0.5$)	Tails ($p = 0.5$)
Take the bet	100	0
Decline	50	50

- **Utilities** are numbers for how good an outcome is, by the agent's lights.
- Where do the numbers come from? Later in the course. For now: **more is better**.
- The **expected utility** of an act: the utilities of its outcomes, weighted by the probabilities of the states.

Dominance

	Miners in A	Miners in B
Block A	10	0
Block B	0	10
Block neither	9	9

Dominance

	Miners in A	Miners in B
Block A	10	0
Block B	0	10
Block neither	9	9
Block both	**10**	**10**

Dominance

	Miners in A	Miners in B
Block A	10	0
Block B	0	10
Block neither	9	9
Block both	**10**	**10**

Weak statewise dominance

Better in **some** state, at least as good in the rest.

Dominance

	Miners in A	Miners in B
Block A	10	0
Block B	0	10
Block neither	9	9
Block both	**10**	**10**

Weak statewise dominance

Better in **some** state, at least as good in the rest.

Strong statewise dominance

Better in **all** states.

Bayesianism: two constraints

Probabilism

Credences must obey the **probability axioms**.

Bayesianism: two constraints

Probabilism

Credences must obey the **probability axioms**.

The Principle of Conditionalisation

They must change in response to evidence by **conditionalisation**.

Bayesianism: two constraints

Probabilism

Credences must obey the **probability axioms**.

The Principle of Conditionalisation

They must change in response to evidence by **conditionalisation**.

Today: both of them.

The sample space

Sample space Ω

All the possible outcomes — all the states of nature you're uncertain about in the situation at hand.

The sample space

Sample space Ω

All the possible outcomes — all the states of nature you're uncertain about in the situation at hand.

- Results of rolling a die. Coin flips. Ways the weather could go.

The sample space

Sample space Ω

All the possible outcomes — all the states of nature you're uncertain about in the situation at hand.

- Results of rolling a die. Coin flips. Ways the weather could go.
- Or the decisions somebody else will make without your input.

The sample space

Sample space Ω

All the possible outcomes — all the states of nature you're uncertain about in the situation at hand.

- Results of rolling a die. Coin flips. Ways the weather could go.
- Or the decisions somebody else will make without your input.

A sample space is just a **set**: a collection of things.

Events

Event

A **set of outcomes**.

Events

Event

A **set of outcomes**.

- The die rolls 2 *or* 3. The weather is rainy *or* overcast.

Events

Event

A **set of outcomes**.

- The die rolls 2 *or* 3. The weather is rainy *or* overcast.
- An event can have just one outcome in it — formally different, but treat them the same.

Events

Event

A **set of outcomes**.

- The die rolls 2 *or* 3. The weather is rainy *or* overcast.
- An event can have just one outcome in it — formally different, but treat them the same.
- The event that *anything at all* happens is Ω itself. Last week I called this the **universal event**.

Events

Event

A **set of outcomes**.

- The die rolls 2 *or* 3. The weather is rainy *or* overcast.
- An event can have just one outcome in it — formally different, but treat them the same.
- The event that *anything at all* happens is Ω itself. Last week I called this the **universal event**.
- The event with **no** outcomes in it is $\varnothing$: the **empty event**. It can never happen.

Combining events

$P \cup Q$ — something in **either** happens

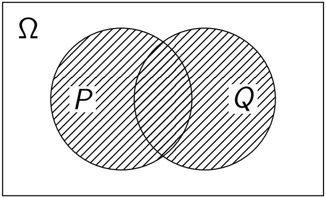

Combining events

$P \cap Q$ — something in **both** happens

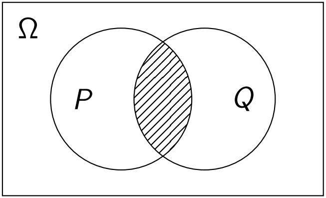

Combining events

P^c — everything in Ω **outside** P

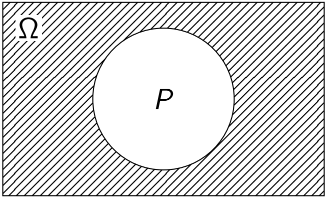

We'll often write $\neg P$ for this.

A note on notation

- Officially: events are **sets**, combined with **union** and **intersection**.

A note on notation

- Officially: events are **sets**, combined with **union** and **intersection**.
- Others — including Schwarz in BDRC — talk about **propositions**, combined with $\wedge$ and $\vee$.

A note on notation

- Officially: events are **sets**, combined with **union** and **intersection**.
- Others — including Schwarz in BDRC — talk about **propositions**, combined with $\wedge$ and $\vee$.
- Strictly these are different things. One takes sets to sets; the other takes truth-apt things to truth-apt things.

A note on notation

- Officially: events are **sets**, combined with **union** and **intersection**.
- Others — including Schwarz in BDRC — talk about **propositions**, combined with $\wedge$ and $\vee$.
- Strictly these are different things. One takes sets to sets; the other takes truth-apt things to truth-apt things.
- But they behave the same way almost everywhere.

A note on notation

- Officially: events are **sets**, combined with **union** and **intersection**.
- Others — including Schwarz in BDRC — talk about **propositions**, combined with $\wedge$ and $\vee$.
- Strictly these are different things. One takes sets to sets; the other takes truth-apt things to truth-apt things.
- But they behave the same way almost everywhere.

I won't penalise you for writing $\wedge$ instead of $\cap$.

One thing I'm going to ignore

- A **probability space** officially comes with a sample space Ω *and* a specification Σ of which subsets count as events.

One thing I'm going to ignore

- A **probability space** officially comes with a sample space Ω *and* a specification Σ of which subsets count as events.
- Why? Because in general **not every subset can coherently be given a probability**, if we want certain intuitively plausible regularities to hold.

One thing I'm going to ignore

- A **probability space** officially comes with a sample space Ω *and* a specification Σ of which subsets count as events.
- Why? Because in general **not every subset can coherently be given a probability**, if we want certain intuitively plausible regularities to hold.
- The reasons are complicated.

One thing I'm going to ignore

- A **probability space** officially comes with a sample space Ω *and* a specification Σ of which subsets count as events.
- Why? Because in general **not every subset can coherently be given a probability**, if we want certain intuitively plausible regularities to hold.
- The reasons are complicated.

We'll assume any set of outcomes is an event.

Mutually exclusive

X and Y have **no outcome in common**: $X \cap Y = \varnothing$

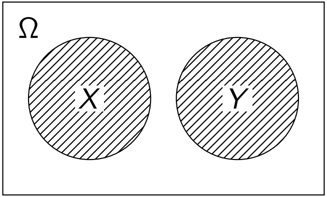

Mutually exclusive

X and Y have **no outcome in common**: $X \cap Y = \varnothing$

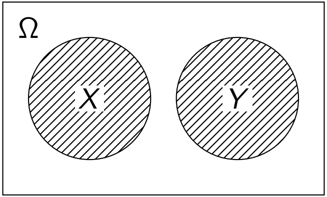

Also called **disjoint**.

Partitions

Jointly exhaustive: $X_1 \cup \cdots \cup X_n = \Omega$

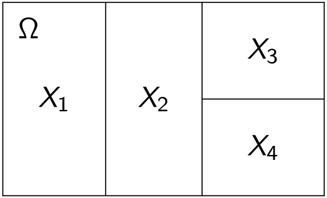

Partitions

Jointly exhaustive: $X_1 \cup \cdots \cup X_n = \Omega$

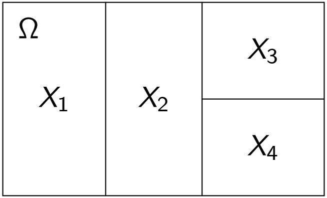

Mutually exclusive **and** jointly exhaustive: a **partition**.

Partitions

Jointly exhaustive: $X_1 \cup \cdots \cup X_n = \Omega$

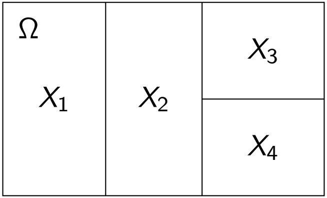

Mutually exclusive **and** jointly exhaustive: a **partition**.

Exactly one of them obtains — like the states of a decision table.

What's wrong with these?

Joe thinks there's a 50% chance it will rain today, and a 20% chance it won't.

What's wrong with these?

Joe thinks there's a 50% chance it will rain today, and a 20% chance it won't.

Biqing thinks there's a 60% chance she failed the exam she just took, and a 70% chance she failed it *badly*.

What's wrong with these?

Joe thinks there's a 50% chance it will rain today, and a 20% chance it won't.

Biqing thinks there's a 60% chance she failed the exam she just took, and a 70% chance she failed it *badly*.

Homer thinks there's a 30% chance he'll drink one beer, a 50% chance he'll drink more than one, and a 70% chance he'll drink at least one.

The axioms — informally, as last week

(i) **Non-negativity.** Probabilities can't be negative.

The axioms — informally, as last week

(i) **Non-negativity.** Probabilities can't be negative.

(ii) **Normalisation.** The probability that *something* happens is 1.

The axioms — informally, as last week

(i) **Non-negativity.** Probabilities can't be negative.
(ii) **Normalisation.** The probability that *something* happens is 1.
(iii) **Additivity.** The probability that *either* of two completely separate events happens is the sum of their probabilities.

The axioms — informally, as last week

(i) **Non-negativity.** Probabilities can't be negative.

(ii) **Normalisation.** The probability that *something* happens is 1.

(iii) **Additivity.** The probability that *either* of two completely separate events happens is the sum of their probabilities.

A die lands 1 *or* 2 with probability $1/6 + 1/6 = 1/3$.

Now we can name the pieces

- Ω is the **sample space**.
- Σ is the **set of events**.
- $\mathrm{Cr}(\sigma)$ is your **credence** in an event $\sigma \in \Sigma$.

Now we can name the pieces

- Ω is the **sample space**.
- Σ is the **set of events**.
- $\mathrm{Cr}(\sigma)$ is your **credence** in an event $\sigma \in \Sigma$.
- "Completely separate" is **disjointness**: $X \cap Y = \varnothing$.

Now we can name the pieces

- Ω is the **sample space**.
- Σ is the **set of events**.
- $Cr(\sigma)$ is your **credence** in an event $\sigma \in \Sigma$.
- "Completely separate" is **disjointness**: $X \cap Y = \varnothing$.

The axioms

(i) $Cr(\sigma) \geq 0$ for every $\sigma \in \Sigma$

Now we can name the pieces

- Ω is the **sample space**.
- Σ is the **set of events**.
- $\mathrm{Cr}(\sigma)$ is your **credence** in an event $\sigma \in \Sigma$.
- "Completely separate" is **disjointness**: $X \cap Y = \varnothing$.

The axioms

(i) $\mathrm{Cr}(\sigma) \geq 0$ for every $\sigma \in \Sigma$

(ii) $\mathrm{Cr}(\Omega) = 1$

Now we can name the pieces

- Ω is the **sample space**.
- Σ is the **set of events**.
- $\mathrm{Cr}(\sigma)$ is your **credence** in an event $\sigma \in \Sigma$.
- "Completely separate" is **disjointness**: $X \cap Y = \varnothing$.

The axioms

(i) $\mathrm{Cr}(\sigma) \geq 0$ for every $\sigma \in \Sigma$

(ii) $\mathrm{Cr}(\Omega) = 1$

(iii) if $X \cap Y = \varnothing$, then $\mathrm{Cr}(X \cup Y) = \mathrm{Cr}(X) + \mathrm{Cr}(Y)$

So what went wrong?

- **Joe.** Rain and no-rain are disjoint and between them exhaust Ω. So they must sum to 1. His sum to 0.7.

So what went wrong?

- **Joe.** Rain and no-rain are disjoint and between them exhaust Ω. So they must sum to 1. His sum to 0.7.
- **Biqing.** Failing badly is a *way of* failing. So it can't be more probable than failing.

So what went wrong?

- **Joe.** Rain and no-rain are disjoint and between them exhaust Ω. So they must sum to 1. His sum to 0.7.
- **Biqing.** Failing badly is a *way of* failing. So it can't be more probable than failing.
- **Homer.** One beer and more than one are disjoint, and together they are "at least one". So that should be $0.3 + 0.5 = 0.8$ — not 0.7.

Conditional probability

Conditional probability

$$\mathrm{Cr}(A \mid B) = \frac{\mathrm{Cr}(A \cap B)}{\mathrm{Cr}(B)} \qquad \text{provided } \; \mathrm{Cr}(B) > 0.$$

Conditional probability

Conditional probability

$$\mathrm{Cr}(A \mid B) = \frac{\mathrm{Cr}(A \cap B)}{\mathrm{Cr}(B)} \qquad \text{provided } \mathrm{Cr}(B) > 0.$$

Your credence in A, **on the supposition that** B.

Conditional probability

Conditional probability

$$\mathrm{Cr}(A \mid B) = \frac{\mathrm{Cr}(A \cap B)}{\mathrm{Cr}(B)} \qquad \text{provided } \ \mathrm{Cr}(B) > 0.$$

Your credence in A, **on the supposition that** B.

Note the proviso: conditioning on something you've ruled out is **undefined**, not zero.

Why that formula?

Before: Ω is everything.

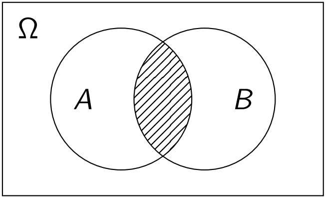

Why that formula?

After: B **is everything.**

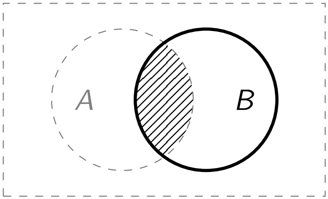

The shaded region is $A \cap B$ both times. What changes is **what it's measured against**.

Probabilistic independence

Probabilistic independence

$$\mathrm{Cr}(A \cap B) = \mathrm{Cr}(A) \cdot \mathrm{Cr}(B)$$

Probabilistic independence

Probabilistic independence

$$\mathrm{Cr}(A \cap B) = \mathrm{Cr}(A) \cdot \mathrm{Cr}(B)$$

Where $\mathrm{Cr}(B) > 0$, that's the same as

$$\mathrm{Cr}(A \mid B) = \mathrm{Cr}(A).$$

Probabilistic independence

Probabilistic independence

$$\mathrm{Cr}(A \cap B) = \mathrm{Cr}(A) \cdot \mathrm{Cr}(B)$$

Where $\mathrm{Cr}(B) > 0$, that's the same as

$$\mathrm{Cr}(A \mid B) = \mathrm{Cr}(A).$$

Learning B tells you **nothing** about A.

Probabilistic independence is not mutual exclusivity

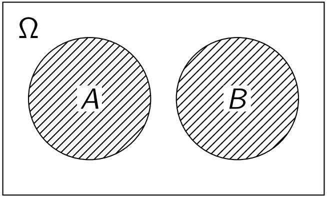

Probabilistic independence is not mutual exclusivity

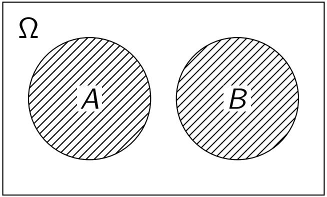

- Disjoint, so $\mathrm{Cr}(A \cap B) = 0$ — but $\mathrm{Cr}(A) \cdot \mathrm{Cr}(B) > 0$.

Probabilistic independence is not mutual exclusivity

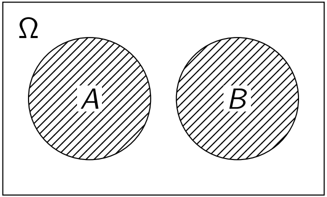

- Disjoint, so $\mathrm{Cr}(A \cap B) = 0$ — but $\mathrm{Cr}(A) \cdot \mathrm{Cr}(B) > 0$.
- So they are **not probabilistically independent**.

Probabilistic independence is not mutual exclusivity

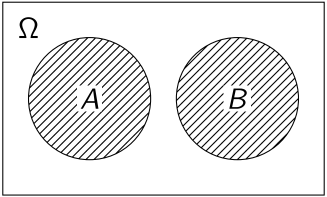

- Disjoint, so $\mathrm{Cr}(A \cap B) = 0$ — but $\mathrm{Cr}(A) \cdot \mathrm{Cr}(B) > 0$.
- So they are **not probabilistically independent**.
- They're as **dependent** as two events can be: learning A takes your credence in B straight to zero.

Conditionalisation

The Principle of Conditionalisation

Suppose your credences are Cr, and you learn E **and nothing more**. Then your new credence in any P should be

$$\mathrm{Cr}_{\mathrm{new}}(P) = \mathrm{Cr}(P \mid E).$$

Conditionalisation

The Principle of Conditionalisation

Suppose your credences are Cr, and you learn E **and nothing more**. Then your new credence in any P should be

$$\mathrm{Cr}_{\mathrm{new}}(P) = \mathrm{Cr}(P \mid E).$$

Your old **conditional** credence becomes your new **unconditional** one.

Conditionalisation, in pictures

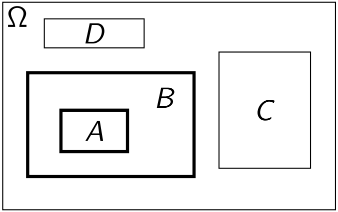

Areas are probabilities. We're going to condition on B.

Conditionalisation, in pictures

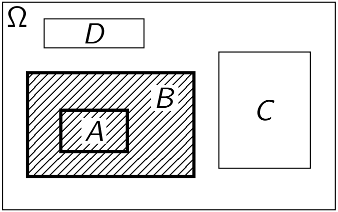

B is the **conditioning event**.

Conditionalisation, in pictures

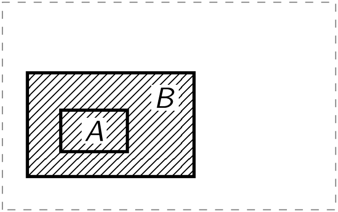

Everything outside B is **discarded**. Nothing has moved yet.

Conditionalisation, in pictures

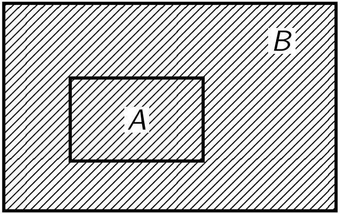

B **blown up** to fill the space Ω used to. Every ratio inside is unchanged.

Conditionalisation, in pictures

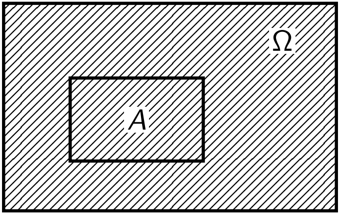

B is the new sample space.

Total probability

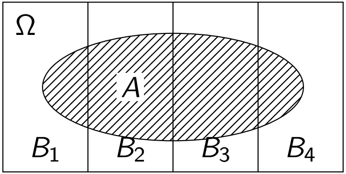

Total probability

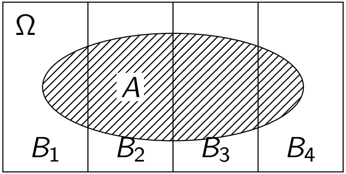

The B_i **partition** Ω, so they cut A into disjoint pieces $A \cap B_i$.

Total probability

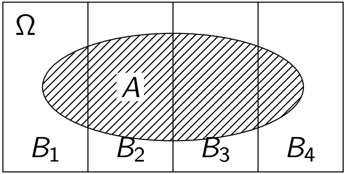

The B_i **partition** Ω, so they cut A into disjoint pieces $A \cap B_i$.

Their probabilities **add up to** $\mathrm{Cr}(A)$.

Bayes' theorem: what we have

$$\mathrm{Cr}(A \mid B) = \frac{\mathrm{Cr}(A \cap B)}{\mathrm{Cr}(B)} \qquad \text{and} \qquad \mathrm{Cr}(B \mid A) = \frac{\mathrm{Cr}(A \cap B)}{\mathrm{Cr}(A)}$$

Bayes' theorem: what we have

$$\mathrm{Cr}(A \mid B) = \frac{\mathrm{Cr}(A \cap B)}{\mathrm{Cr}(B)} \qquad \text{and} \qquad \mathrm{Cr}(B \mid A) = \frac{\mathrm{Cr}(A \cap B)}{\mathrm{Cr}(A)}$$

What we want: a way across, from one to the other.

Bayes' theorem: what we have

$$\mathrm{Cr}(A \mid B) = \frac{\mathrm{Cr}(A \cap B)}{\mathrm{Cr}(B)} \qquad \text{and} \qquad \mathrm{Cr}(B \mid A) = \frac{\mathrm{Cr}(A \cap B)}{\mathrm{Cr}(A)}$$

What we want: a way across, from one to the other.

Bayes' theorem

$$\mathrm{Cr}(A \mid B) = \frac{\mathrm{Cr}(B \mid A)\ \mathrm{Cr}(A)}{\mathrm{Cr}(B)} \qquad \text{provided } \mathrm{Cr}(A), \mathrm{Cr}(B) > 0.$$

Bayes factors

Two hypotheses H_1, H_2, and some evidence E.

Bayes factors

Two hypotheses H_1, H_2, and some evidence E.

Bayes factor

$$\frac{\mathrm{Cr}(E \mid H_1)}{\mathrm{Cr}(E \mid H_2)}$$

Bayes factors

Two hypotheses H_1, H_2, and some evidence E.

Bayes factor

$$\frac{Cr(E \mid H_1)}{Cr(E \mid H_2)}$$

$$\underbrace{\frac{Cr(H_1 \mid E)}{Cr(H_2 \mid E)}}_{\text{posterior odds}} = \underbrace{\frac{Cr(E \mid H_1)}{Cr(E \mid H_2)}}_{\text{Bayes factor}} \times \underbrace{\frac{Cr(H_1)}{Cr(H_2)}}_{\text{prior odds}}$$

Bayes factors

Two hypotheses H_1, H_2, and some evidence E.

Bayes factor

$$\frac{\mathrm{Cr}(E \mid H_1)}{\mathrm{Cr}(E \mid H_2)}$$

$$\underbrace{\frac{\mathrm{Cr}(H_1 \mid E)}{\mathrm{Cr}(H_2 \mid E)}}_{\text{posterior odds}} = \underbrace{\frac{\mathrm{Cr}(E \mid H_1)}{\mathrm{Cr}(E \mid H_2)}}_{\text{Bayes factor}} \times \underbrace{\frac{\mathrm{Cr}(H_1)}{\mathrm{Cr}(H_2)}}_{\text{prior odds}}$$

The evidence supplies the **multiplier**. Your prior says **where you started**.

Bayes factors

Two hypotheses H_1, H_2, and some evidence E.

Bayes factor

$$\frac{\mathrm{Cr}(E \mid H_1)}{\mathrm{Cr}(E \mid H_2)}$$

$$\underbrace{\frac{\mathrm{Cr}(H_1 \mid E)}{\mathrm{Cr}(H_2 \mid E)}}_{\text{posterior odds}} = \underbrace{\frac{\mathrm{Cr}(E \mid H_1)}{\mathrm{Cr}(E \mid H_2)}}_{\text{Bayes factor}} \times \underbrace{\frac{\mathrm{Cr}(H_1)}{\mathrm{Cr}(H_2)}}_{\text{prior odds}}$$

The evidence supplies the **multiplier**. Your prior says **where you started**.

A Bayes factor of 1 moves you not at all.

What a Dutch book argument is

The best-known arguments for **probabilism** — for the three axioms we saw earlier.

What a Dutch book argument is

The best-known arguments for **probabilism** — for the three axioms we saw earlier.

- They aim to show that agents whose credences break the axioms are **systematically exploitable**.

What a Dutch book argument is

The best-known arguments for **probabilism** — for the three axioms we saw earlier.

- They aim to show that agents whose credences break the axioms are **systematically exploitable**.
- They can end up taking actions which together amount to **choosing to give up money** — or something else they value — **for no gain whatsoever**.

What a Dutch book argument is

The best-known arguments for **probabilism** — for the three axioms we saw earlier.

- They aim to show that agents whose credences break the axioms are **systematically exploitable**.
- They can end up taking actions which together amount to **choosing to give up money** — or something else they value — **for no gain whatsoever**.
- And being systematically exploitable is irrational. So rationality requires you not to have credences like that.

What a Dutch book argument is

The best-known arguments for **probabilism** — for the three axioms we saw earlier.

- They aim to show that agents whose credences break the axioms are **systematically exploitable**.
- They can end up taking actions which together amount to **choosing to give up money** — or something else they value — **for no gain whatsoever**.
- And being systematically exploitable is irrational. So rationality requires you not to have credences like that.

There's also a Dutch book argument for **conditionalisation**. Not this week — maybe week 5.

Linking belief to action

Write a **\$100 bet on** A for a deal paying \$100 if A is true, nothing otherwise.

Linking belief to action

Write a **\$100 bet on** A for a deal paying \$100 if A is true, nothing otherwise.

The betting interpretation

An agent whose credence in A is p will **buy** a \$100 bet on A for any price below \$100p, and **sell** one for any price above it.

Linking belief to action

Write a **\$100 bet on** A for a deal paying \$100 if A is true, nothing otherwise.

The betting interpretation

An agent whose credence in A is p will **buy** a \$100 bet on A for any price below \$100p, and **sell** one for any price above it.

Her credence fixes a price, and she'll take **either side** of it.

Linking belief to action

Write a **\$100 bet on** A for a deal paying \$100 if A is true, nothing otherwise.

The betting interpretation

An agent whose credence in A is p will **buy** a \$100 bet on A for any price below \$100p, and **sell** one for any price above it.

Her credence fixes a price, and she'll take **either side** of it.

We'll assume this today. It's a big assumption. Complaints held over.

Non-negativity, in one bet

Suppose $\mathrm{Cr}(A) = -0.5$.

Non-negativity, in one bet

Suppose $\mathrm{Cr}(A) = -0.5$.

- She'll sell a \$100 bet on A for any price above $-\$50$ — and a **negative price means she pays**.

Non-negativity, in one bet

Suppose $\mathrm{Cr}(A) = -0.5$.

- She'll sell a \$100 bet on A for any price above $-\$50$ — and a **negative price means she pays**.
- So she'll happily pay someone \$49 to take the bet on — and still owes \$100 if A.

Non-negativity, in one bet

Suppose $\mathrm{Cr}(A) = -0.5$.

- She'll sell a \$100 bet on A for any price above $-\$50$ — and a **negative price means she pays**.
- So she'll happily pay someone \$49 to take the bet on — and still owes \$100 if A.

	A	$\neg A$
Her net gain	$-\$149$	$-\$49$

Non-negativity, in one bet

Suppose $\mathrm{Cr}(A) = -0.5$.

- She'll sell a \$100 bet on A for any price above $-\$50$ — and a **negative price means she pays**.
- So she'll happily pay someone \$49 to take the bet on — and still owes \$100 if A.

	A	$\neg A$
Her net gain	$-\$149$	$-\$49$

She loses \$49 **whatever happens.**

Normalisation, in one bet

Suppose $\mathrm{Cr}(\Omega) = 0.9$.

Normalisation, in one bet

Suppose $\text{Cr}(\Omega) = 0.9$.

- She'll sell a \$100 bet on Ω for \$91.

Normalisation, in one bet

Suppose $\mathrm{Cr}(\Omega) = 0.9$.

- She'll sell a \$100 bet on Ω for \$91.
- But Ω is **guaranteed**. She collects \$91 and pays out \$100.

Normalisation, in one bet

Suppose $\mathrm{Cr}(\Omega) = 0.9$.

- She'll sell a \$100 bet on Ω for \$91.
- But Ω is **guaranteed**. She collects \$91 and pays out \$100.
- Net: $-\$9$, for certain.

Normalisation, in one bet

Suppose $Cr(\Omega) = 0.9$.

- She'll sell a \$100 bet on Ω for \$91.
- But Ω is **guaranteed**. She collects \$91 and pays out \$100.
- Net: $-\$9$, for certain.

And if $Cr(\Omega) = 1.1$? She buys the same bet for \$109 and collects \$100. **Same \$9 loss.**

Finite additivity is harder

These worked because **one bet** did the damage.

Finite additivity is harder

These worked because **one bet** did the damage.

Finite additivity is about how credences relate to **other credences**.

Finite additivity is harder

These worked because **one bet** did the damage.

Finite additivity is about how credences relate to **other credences**.

So we'll need **more than one bet**.

A violation of finite additivity

P and Q are **mutually exclusive**. Our agent has credence **0.2** in P, **0.2** in Q, but **0.8** in $P \cup Q$ — where additivity requires **0.4**.

A violation of finite additivity

P and Q are **mutually exclusive**. Our agent has credence **0.2** in P, **0.2** in Q, but **0.8** in $P \cup Q$ — where additivity requires **0.4**.

Bet 1 She **sells** a \$100 bet on P for \$21. Her price is \$20, so selling above it is a gain.

A violation of finite additivity

P and Q are **mutually exclusive**. Our agent has credence **0.2** in P, **0.2** in Q, but **0.8** in $P \cup Q$ — where additivity requires **0.4**.

Bet 1 She **sells** a \$100 bet on P for \$21. Her price is \$20, so selling above it is a gain.

Bet 2 She **sells** a \$100 bet on Q for \$21. Likewise.

A violation of finite additivity

P and Q are **mutually exclusive**. Our agent has credence **0.2** in P, **0.2** in Q, but **0.8** in $P \cup Q$ — where additivity requires **0.4**.

Bet 1 She **sells** a \$100 bet on P for \$21. Her price is \$20, so selling above it is a gain.

Bet 2 She **sells** a \$100 bet on Q for \$21. Likewise.

Bet 3 She **buys** a \$100 bet on $P \cup Q$ for \$79. Her price is \$80, so buying below it is a gain.

A violation of finite additivity

P and Q are **mutually exclusive**. Our agent has credence **0.2** in P, **0.2** in Q, but **0.8** in $P \cup Q$ — where additivity requires **0.4**.

Bet 1 She **sells** a \$100 bet on P for \$21. Her price is \$20, so selling above it is a gain.

Bet 2 She **sells** a \$100 bet on Q for \$21. Likewise.

Bet 3 She **buys** a \$100 bet on $P \cup Q$ for \$79. Her price is \$80, so buying below it is a gain.

Each one, on its own, is a bet she **wants**.

But take all three. . .

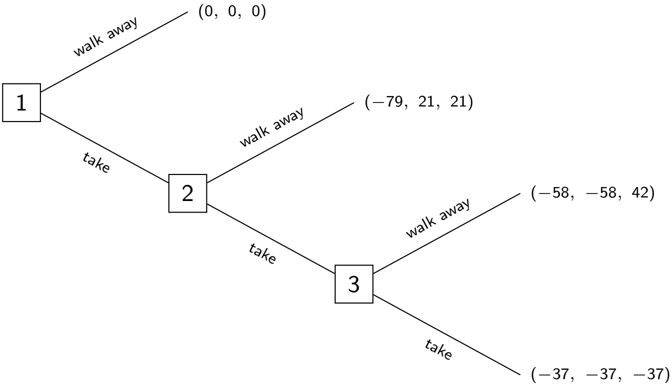

But take all three. . .

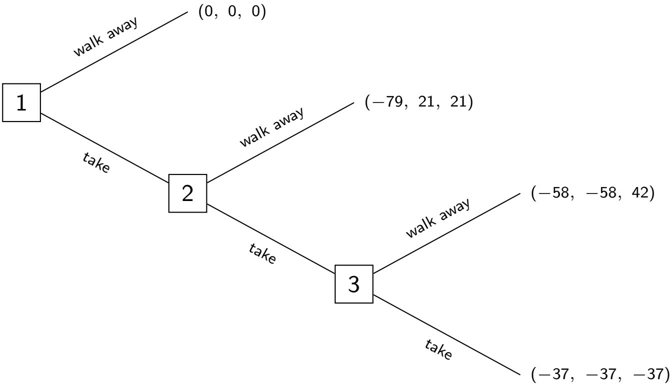

She is down **$37**, however things turn out.

But she can see it coming

- So far we've assumed she treats each offer as if it were **the only decision she'll ever face**. That's **myopic** choice.

But she can see it coming

- So far we've assumed she treats each offer as if it were **the only decision she'll ever face**. That's **myopic** choice.
- An agent who **looks ahead**, and takes into account what she'll do later, is **sophisticated**.

But she can see it coming

- So far we've assumed she treats each offer as if it were **the only decision she'll ever face**. That's **myopic** choice.
- An agent who **looks ahead**, and takes into account what she'll do later, is **sophisticated**.

Suppose she did take all three. Then she'd go on taking bets at node 2 too. So at node 1 her choice is really between

But she can see it coming

- So far we've assumed she treats each offer as if it were **the only decision she'll ever face**. That's **myopic** choice.
- An agent who **looks ahead**, and takes into account what she'll do later, is **sophisticated**.

Suppose she did take all three. Then she'd go on taking bets at node 2 too. So at node 1 her choice is really between

$$(0,\ 0,\ 0) \qquad \text{and} \qquad (-37,\ -37,\ -37)$$

But she can see it coming

- So far we've assumed she treats each offer as if it were **the only decision she'll ever face**. That's **myopic** choice.
- An agent who **looks ahead**, and takes into account what she'll do later, is **sophisticated**.

Suppose she did take all three. Then she'd go on taking bets at node 2 too. So at node 1 her choice is really between

$$(0,\ 0,\ 0) \qquad \text{and} \qquad (-37,\ -37,\ -37)$$

So she walks away at node 1.

Can the book be repaired?

It can.

Can the book be repaired?

It can.

The Dutch book for finite additivity can be rebuilt so that it works even against an agent using sophisticated choice — which is what she **ought** to be using.

Can the book be repaired?

It can.

The Dutch book for finite additivity can be rebuilt so that it works even against an agent using sophisticated choice — which is what she **ought** to be using.

Week 5, if there's time.

From the theorem to an argument

The flat-footed version: *a cunning Dutchman might come along and take your money, so you'd better have probabilistic credences.*

From the theorem to an argument

The flat-footed version: *a cunning Dutchman might come along and take your money, so you'd better have probabilistic credences.*

- On this version, the reason to avoid non-probabilistic credences is that it **protects you from actually being exploited**.

From the theorem to an argument

The flat-footed version: *a cunning Dutchman might come along and take your money, so you'd better have probabilistic credences.*

- On this version, the reason to avoid non-probabilistic credences is that it **protects you from actually being exploited**.
- But suppose you knew for certain that **no Dutch book situation would arise**. Then that reason vanishes.

From the theorem to an argument

The flat-footed version: *a cunning Dutchman might come along and take your money, so you'd better have probabilistic credences.*

- On this version, the reason to avoid non-probabilistic credences is that it **protects you from actually being exploited**.
- But suppose you knew for certain that **no Dutch book situation would arise**. Then that reason vanishes.
- A proponent of probabilism wouldn't call you rational. It's the **pattern of credences itself** that was supposed to be irrational.

A better version

1. If you **do not** bet in line with the betting interpretation for the credences you have, you are irrational.

A better version

1. If you **do not** bet in line with the betting interpretation for the credences you have, you are irrational.
2. If you have non-probabilistic credences **and** bet in line with the betting interpretation, you are vulnerable to a Dutch book.

A better version

1. If you **do not** bet in line with the betting interpretation for the credences you have, you are irrational.
2. If you have non-probabilistic credences **and** bet in line with the betting interpretation, you are vulnerable to a Dutch book.
3. If you are vulnerable to a Dutch book, **you** are irrational.

A better version

1. If you **do not** bet in line with the betting interpretation for the credences you have, you are irrational.
2. If you have non-probabilistic credences **and** bet in line with the betting interpretation, you are vulnerable to a Dutch book.
3. If you are vulnerable to a Dutch book, **you** are irrational.

Either you bet in line with it or you don't. By (1) if you don't; by (2) and (3) if you do.

A better version

1. If you **do not** bet in line with the betting interpretation for the credences you have, you are irrational.
2. If you have non-probabilistic credences **and** bet in line with the betting interpretation, you are vulnerable to a Dutch book.
3. If you are vulnerable to a Dutch book, **you** are irrational.

Either you bet in line with it or you don't. By (1) if you don't; by (2) and (3) if you do.

So: if your credences are non-probabilistic, **you** are irrational.

And it isn't about money

You're not rationally required to value money.

And it isn't about money

You're not rationally required to value money.

The argument works for **any** quantity you do care about — so long as you're required to bet in line with the betting interpretation for it.

For next week

Newcomb's problem

BDRC ch. 9, and Adam Elga, "Newcomb University: A Play in One Act", *Analysis* 80.2 (2020), pp. 212–221.

For next week

Newcomb's problem

BDRC ch. 9, and Adam Elga, "Newcomb University: A Play in One Act", *Analysis* 80.2 (2020), pp. 212–221.

See you next week.