Decision Theory and Social Choice

PE3101P — Lecture 1

Tomi Francis

National University of Singapore

13^{th} August, 2026

Today

Hello

- I'm your instructor for **PE3101P: Decision Theory and Social Choice**.

Hello

- I'm your instructor for **PE3101P: Decision Theory and Social Choice**.
- Please call me **Tomi**.

Hello

- I'm your instructor for **PE3101P: Decision Theory and Social Choice**.
- Please call me **Tomi**.
- No need for "Dr. Francis" or "Prof. Francis" or anything like that.

The register

- The University wants attendance records.

The register

- The University wants attendance records.
- **Please sign the register** as it comes round.

The register

- The University wants attendance records.
- **Please sign the register** as it comes round.
- Add your name and signature. In future we'll have printed copies.

Finding me

- **Email:** tomi.francis@nus.edu.sg

Finding me

- **Email:** tomi.francis@nus.edu.sg
- **Office hours:** Thursdays, 4–6pm, **AS7-03-12**.

Finding me

- **Email:** tomi.francis@nus.edu.sg
- **Office hours:** Thursdays, 4–6pm, **AS7-03-12**.
- That's straight after this class — no need to wait until 4.

Finding me

- **Email:** tomi.francis@nus.edu.sg
- **Office hours:** Thursdays, 4–6pm, **AS7-03-12**.
- That's straight after this class — no need to wait until 4.
- Need me outside those hours? Email me and we'll arrange it.

Finding me

- **Email:** tomi.francis@nus.edu.sg
- **Office hours:** Thursdays, 4–6pm, **AS7-03-12**.
- That's straight after this class — no need to wait until 4.
- Need me outside those hours? Email me and we'll arrange it.
- (My office may move later this semester. I'll tell you if it does.)

What this course is about

Decision-making in two kinds of case:

What this course is about

Decision-making in two kinds of case:

(1) We don't know how our actions will turn out, and must weigh the risks.

What this course is about

Decision-making in two kinds of case:

(1) We don't know how our actions will turn out, and must weigh the risks.

(2) Different people want different things, and we must decide for everybody.

What this course is about

Decision-making in two kinds of case:

(1) We don't know how our actions will turn out, and must weigh the risks.

(2) Different people want different things, and we must decide for everybody.

- (1) is **decision theory**.
- (2) is **social choice theory**.

What this course is about

Decision-making in two kinds of case:

(1) We don't know how our actions will turn out, and must weigh the risks.

(2) Different people want different things, and we must decide for everybody.

- (1) is **decision theory**.
- (2) is **social choice theory**.

Roughly half the semester on each. Decision theory first.

Where the course materials live

- **My website is the primary source.**

Where the course materials live

- **My website is the primary source.**
- Canvas when I can, but the website will be better and more current.

Where the course materials live

- **My website is the primary source.**
- Canvas when I can, but the website will be better and more current.
- New material each week — lecture notes, optional assignments, solutions.

Where the course materials live

- **My website is the primary source.**
- Canvas when I can, but the website will be better and more current.
- New material each week — lecture notes, optional assignments, solutions.
- **Usually up by the end of Friday**, the day after the lecture.

Where the course materials live

- **My website is the primary source.**
- Canvas when I can, but the website will be better and more current.
- New material each week — lecture notes, optional assignments, solutions.
- **Usually up by the end of Friday**, the day after the lecture.
- Today's notes are already there.

Where the course materials live

- **My website is the primary source.**
- Canvas when I can, but the website will be better and more current.
- New material each week — lecture notes, optional assignments, solutions.
- **Usually up by the end of Friday**, the day after the lecture.
- Today's notes are already there.
- Trouble getting at anything? Email me.

The syllabus

- A printed copy should be in front of you — if not, grab one from the front.

The syllabus

- A printed copy should be in front of you — if not, grab one from the front.
- Also on my website. (And on Canvas when I can, but the website is likelier to be current.)

The syllabus

- A printed copy should be in front of you — if not, grab one from the front.
- Also on my website. (And on Canvas when I can, but the website is likelier to be current.)
- **Logistical question? Reread the syllabus first.** It may well answer it.

The textbook

Schwarz, *Belief, Desire, and Rational Choice* (2022)

Open access. A PDF is on my website.

The textbook

Schwarz, *Belief, Desire, and Rational Choice* (2022)

Open access. A PDF is on my website.

- We'll follow it loosely.

The textbook

Schwarz, *Belief, Desire, and Rational Choice* (2022)

Open access. A PDF is on my website.

- We'll follow it loosely.
- **Please use the copy on my website**, so we're all reading the same one.

Readings

- Usually **1–2 short readings** per week.

Readings

- Usually **1–2 short readings** per week.
- Chosen for (1) informativeness and (2) fun. There are some real crackers.

Readings

- Usually **1–2 short readings** per week.
- Chosen for (1) informativeness and (2) fun. There are some real crackers.
- **Read them before class.** The lectures will be easier to follow.

Readings

- Usually **1–2 short readings** per week.
- Chosen for (1) informativeness and (2) fun. There are some real crackers.
- **Read them before class.** The lectures will be easier to follow.
- Not this week, obviously.

Readings

- Usually **1–2 short readings** per week.
- Chosen for (1) informativeness and (2) fun. There are some real crackers.
- **Read them before class.** The lectures will be easier to follow.
- Not this week, obviously.
- Listed in the syllabus — see my website.

Questions and discussion

- Questions are welcome **at any point**. Just raise your hand.

Questions and discussion

- Questions are welcome **at any point**. Just raise your hand.
- I may occasionally decline if I'm rushing to finish something.

Questions and discussion

- Questions are welcome **at any point**. Just raise your hand.
- I may occasionally decline if I'm rushing to finish something.
- **That is not a reason to stop raising your hand.**

Questions and discussion

- Questions are welcome **at any point**. Just raise your hand.
- I may occasionally decline if I'm rushing to finish something.
- **That is not a reason to stop raising your hand.**
- Tutorials are more discussion-focused — I'll want more from you there.

Use of AI

Assessed work

Do not use AI in any way. This is cheating.

Use of AI

Assessed work

Do not use AI in any way. This is cheating.

Everything else

Use it or don't — but use your judgement. Help your learning, don't inhibit it.

Use of AI: good ideas

- Generating extra exercises based on the optional assignments.

Use of AI: good ideas

- Generating extra exercises based on the optional assignments.
- Discussing the readings *after* you've read them yourself.

Use of AI: good ideas

- Generating extra exercises based on the optional assignments.
- Discussing the readings *after* you've read them yourself.
- Building lists of question types and key facts for revision.

Use of AI: good ideas

- Generating extra exercises based on the optional assignments.
- Discussing the readings *after* you've read them yourself.
- Building lists of question types and key facts for revision.
- Asking for explanations of things you find difficult.

Use of AI: good ideas

- Generating extra exercises based on the optional assignments.
- Discussing the readings *after* you've read them yourself.
- Building lists of question types and key facts for revision.
- Asking for explanations of things you find difficult.
- Triaging extra readings you've found, to see if they're worth your time.

Use of AI: bad ideas

- Getting a chatbot to do an optional assignment for you.

Use of AI: bad ideas

- Getting a chatbot to do an optional assignment for you.
- Asking for the answer before spending **ten minutes** on it yourself.

Use of AI: bad ideas

- Getting a chatbot to do an optional assignment for you.
- Asking for the answer before spending **ten minutes** on it yourself.
- Reading AI summaries instead of the readings.

Use of AI: bad ideas

- Getting a chatbot to do an optional assignment for you.
- Asking for the answer before spending **ten minutes** on it yourself.
- Reading AI summaries instead of the readings.
- **Anything that involves switching off your brain.**

Use of AI: bad ideas

- Getting a chatbot to do an optional assignment for you.
- Asking for the answer before spending **ten minutes** on it yourself.
- Reading AI summaries instead of the readings.
- **Anything that involves switching off your brain.**

In short

Do use AI to supplement your learning. **Do not** use it to replace your learning.

We'll start with tables like this

	It rains	It doesn't rain
Take the umbrella	Dry, but carrying it	Dry, but carrying it
Leave the umbrella	Soaked	Dry, hands free

We'll start with tables like this

	It rains	It doesn't rain
Take the umbrella	Dry, but carrying it	Dry, but carrying it
Leave the umbrella	Soaked	Dry, hands free

- We'll learn to pick out the **states**, the **acts** and the **outcomes**.

We'll start with tables like this

	It rains	It doesn't rain
Take the umbrella	Dry, but carrying it	Dry, but carrying it
Leave the umbrella	Soaked	Dry, hands free

- We'll learn to pick out the **states**, the **acts** and the **outcomes**.
- Then **dominance principles** — and why we must be careful with them.

Decision theory: what's coming

- **Probabilities** for states of nature: what they mean, what rules they obey, and why it's irrational to break those rules.

Decision theory: what's coming

- **Probabilities** for states of nature: what they mean, what rules they obey, and why it's irrational to break those rules.
- **Newcomb's problem**: what to do when you can **expect** things to go badly if you perform a certain act, even though you **know** performing that act will make things better.

Decision theory: what's coming

- **Probabilities** for states of nature: what they mean, what rules they obey, and why it's irrational to break those rules.
- **Newcomb's problem**: what to do when you can **expect** things to go badly if you perform a certain act, even though you **know** performing that act will make things better.
- (If that sounds impossible — stick around.)

Decision theory: the centrepiece

Expected Utility Theory

If you choose under risk in certain ways — ways many think rationality requires — then you **must** be acting *as though* you were maximising expected utility.

Decision theory: the centrepiece

Expected Utility Theory

If you choose under risk in certain ways — ways many think rationality requires — then you **must** be acting *as though* you were maximising expected utility.

- i.e. as though you'd put numbers on how good each outcome is, and were making the average as high as possible.

Decision theory: the centrepiece

Expected Utility Theory

If you choose under risk in certain ways — ways many think rationality requires — then you **must** be acting *as though* you were maximising expected utility.

- i.e. as though you'd put numbers on how good each outcome is, and were making the average as high as possible.
- Then: **money pump arguments** — if you don't choose that way, you can be made worse off for no reason, **and you can see it coming**.

Social choice theory: what's coming

- **Harsanyi**: add up everyone's utilities and pick the biggest total. How the argument works — and what's wrong with it.

Social choice theory: what's coming

- **Harsanyi**: add up everyone's utilities and pick the biggest total. How the argument works — and what's wrong with it.
- **Game theory**: when the outcome depends on many people, each anticipating the others.

Social choice theory: what's coming

- **Harsanyi**: add up everyone's utilities and pick the biggest total. How the argument works — and what's wrong with it.
- **Game theory**: when the outcome depends on many people, each anticipating the others.
- **Sen's paradox**: even a minimal commitment to making people better off collides with liberal rights.

Social choice theory: what's coming

- **Harsanyi**: add up everyone's utilities and pick the biggest total. How the argument works — and what's wrong with it.
- **Game theory**: when the outcome depends on many people, each anticipating the others.
- **Sen's paradox**: even a minimal commitment to making people better off collides with liberal rights.
- **Arrow**: the impossibility of a satisfactory voting system.

Social choice theory: what's coming

- **Harsanyi**: add up everyone's utilities and pick the biggest total. How the argument works — and what's wrong with it.
- **Game theory**: when the outcome depends on many people, each anticipating the others.
- **Sen's paradox**: even a minimal commitment to making people better off collides with liberal rights.
- **Arrow**: the impossibility of a satisfactory voting system.
- **Is voting irrational?** You almost certainly won't change the result.

So — let's start on decision theory.

The Miners Problem

The Miners Problem

Ten miners are trapped in a shaft, threatened by rising water. You don't know whether they are in shaft A or shaft B.

The Miners Problem

The Miners Problem

Ten miners are trapped in a shaft, threatened by rising water. You don't know whether they are in shaft A or shaft B. You have enough sandbags to block **one** shaft, not both.

The Miners Problem

The Miners Problem

Ten miners are trapped in a shaft, threatened by rising water. You don't know whether they are in shaft A or shaft B. You have enough sandbags to block **one** shaft, not both. Block the right shaft and all survive; block the wrong one and all die.

The Miners Problem

The Miners Problem

Ten miners are trapped in a shaft, threatened by rising water. You don't know whether they are in shaft A or shaft B. You have enough sandbags to block **one** shaft, not both. Block the right shaft and all survive; block the wrong one and all die. Do nothing, and both shafts half-fill: **one miner dies**.

The Miners Problem

	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

The Miners Problem

	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

What should you do?

Three questions

- Should you block A, block B, or neither?

Three questions

- Should you block A, block B, or neither?
- What if the miners *are* in shaft A? What then?

Three questions

- Should you block A, block B, or neither?
- What if the miners *are* in shaft A? What then?
- What if the miners are in shaft A **and you still don't know it**?

Two senses of "ought"

- If they *are* in A, then in one sense the right thing is to block A — whether or not you know it.

Two senses of "ought"

- If they *are* in A, then in one sense the right thing is to block A — whether or not you know it.
- In another sense, blocking A is reckless, even if you get lucky.

Two senses of "ought"

- If they *are* in A, then in one sense the right thing is to block A — whether or not you know it.
- In another sense, blocking A is reckless, even if you get lucky.

Fact-relative

What you ought to do is whatever will *in fact* be best. You may have no idea what that is. Never mind.

Two senses of "ought"

- If they *are* in A, then in one sense the right thing is to block A — whether or not you know it.
- In another sense, blocking A is reckless, even if you get lucky.

Fact-relative

What you ought to do is whatever will *in fact* be best. You may have no idea what that is. Never mind.

Evidence-relative

What you ought to do is whatever is best *given your evidence*. Here, what you know matters a lot.

Which one is this course about?

Decision theory is about the evidence-relative sense.

Which one is this course about?

Decision theory is about the evidence-relative sense.

How can we decide well, on the basis of the evidence we actually have?

Reading the table

	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

Reading the table

	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

The **columns** are the **states** — or *states of nature*.

- Ways the world might be that you're **uncertain** about.
- That you **can't change**.
- That **determine** what each act leads to.

Reading the table

	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

The **rows** are the **acts**.

- The choices open to you.
- Together with the state, they fix what happens.

Reading the table

	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**

The **cells** are the **outcomes** — or *final outcomes*, or *possible worlds*.

- A complete description of a possible result of your decision.

A question

	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**

Is the result of **Block neither** when the miners are in A **the same outcome** as when they're in B?

Suppose we had a fourth option

	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

Suppose we had a fourth option

	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
Block both	**all 10 live**	**all 10 live**

Suppose we had a fourth option

	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
Block both	**all 10 live**	**all 10 live**

- vs **Block** A: better if they're in B, same if in A.

Suppose we had a fourth option

	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
Block both	**all 10 live**	**all 10 live**

- vs **Block** A: better if they're in B, same if in A.
- vs **Block** B: better if they're in A, same if in B.

Suppose we had a fourth option

	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
Block both	**all 10 live**	**all 10 live**

- vs **Block** A: better if they're in B, same if in A.
- vs **Block** B: better if they're in A, same if in B.
- vs **Block neither**: better **whichever** state obtains.

Two kinds of dominance

Weak statewise dominance

Definitely at least as good, and **possibly** better.

Two kinds of dominance

Weak statewise dominance

Definitely at least as good, and **possibly** better.

So **Block both** weakly statewise-dominates blocking A alone, and blocking B alone.

Two kinds of dominance

Weak statewise dominance

Definitely at least as good, and **possibly** better.

So **Block both** weakly statewise-dominates blocking A alone, and blocking B alone.

Strong statewise dominance

Better **no matter which state obtains**.

Two kinds of dominance

Weak statewise dominance

Definitely at least as good, and **possibly** better.

So **Block both** weakly statewise-dominates blocking A alone, and blocking B alone.

Strong statewise dominance

Better **no matter which state obtains**.

So **Block both** strongly statewise-dominates **Block neither**.

So should we always choose dominant acts?

It seems so.

So should we always choose dominant acts?

It seems so.

Why accept a risk — or a certainty — of a worse outcome?

A counterexample?

A student wonders whether to study for an exam.

	Will Pass	Won't Pass
Study	Pass & No Fun	Fail & No Fun
Don't Study	Pass & Fun	Fail & Fun

A counterexample?

A student wonders whether to study for an exam.

	Will Pass	Won't Pass
Study	Pass & No Fun	Fail & No Fun
Don't Study	Pass & Fun	Fail & Fun

Does "Don't study" statewise dominate "Study"?

Yes — and that's the problem

- The lesson is **not** that statewise dominance is wrong.

Yes — and that's the problem

- The lesson is **not** that statewise dominance is wrong.
- It's that we set up the **states** badly.

Yes — and that's the problem

- The lesson is **not** that statewise dominance is wrong.
- It's that we set up the **states** badly.

Remember

States of nature are things you **can't change**.

Yes — and that's the problem

- The lesson is **not** that statewise dominance is wrong.
- It's that we set up the **states** badly.

Remember

States of nature are things you **can't change**.

But studying *does* make a difference to whether you pass.

The fix

- We should choose statewise-dominant acts — **but only once the states are set up properly**.

The fix

- We should choose statewise-dominant acts — **but only once the states are set up properly**.
- First approximation: the states must be **probabilistically independent** of the act chosen.

The fix

- We should choose statewise-dominant acts — **but only once the states are set up properly**.
- First approximation: the states must be **probabilistically independent** of the act chosen.
- (We'll see what that means later.)

Probabilities matter too

What we should do depends not just on the states and outcomes, but on **how likely** each state is.

Probabilities matter too

What we should do depends not just on the states and outcomes, but on **how likely** each state is.

	World ends tomorrow	It doesn't
Wild party, on credit	A magnificent last night	Years of debt
Stay home, play Scrabble	A quiet last night	A quiet night, no d

Probabilities matter too

What we should do depends not just on the states and outcomes, but on **how likely** each state is.

	World ends tomorrow	It doesn't
Wild party, on credit	A magnificent last night	Years of debt
Stay home, play Scrabble	A quiet last night	A quiet night, no d

- If the world were very likely to end — have the party.

Probabilities matter too

What we should do depends not just on the states and outcomes, but on **how likely** each state is.

	World ends tomorrow	It doesn't
Wild party, on credit	A magnificent last night	Years of debt
Stay home, play Scrabble	A quiet last night	A quiet night, no d

- If the world were very likely to end — have the party.
- In the actual world — Scrabble wins.

Credences

- The probabilities we care about are **credences**, or *degrees of belief*.

Credences

- The probabilities we care about are **credences**, or *degrees of belief*.
- Roughly: how likely **you think** something is, given your evidence.

Credences

- The probabilities we care about are **credences**, or *degrees of belief*.
- Roughly: how likely **you think** something is, given your evidence.
- When I say "probability", I mean **credence**, unless I say otherwise.

Credences aren't objective chances

Jack and Jill

Jack has flipped a coin and is looking at it. It landed heads.

Credences aren't objective chances

Jack and Jill

Jack has flipped a coin and is looking at it. It landed heads. Jill didn't see it, and is trying to guess.

Credences aren't objective chances

Jack and Jill

Jack has flipped a coin and is looking at it. It landed heads. Jill didn't see it, and is trying to guess.

- The **chance** of tails is **zero**. It landed heads.

Credences aren't objective chances

Jack and Jill

Jack has flipped a coin and is looking at it. It landed heads. Jill didn't see it, and is trying to guess.

- The **chance** of tails is **zero**. It landed heads.
- Jill's **credence** in tails should be **50%**.

Credences aren't objective chances

Jack and Jill

Jack has flipped a coin and is looking at it. It landed heads. Jill didn't see it, and is trying to guess.

- The **chance** of tails is **zero**. It landed heads.
- Jill's **credence** in tails should be **50%**.

We care about credences because we often don't know the chances — and we want to decide well on the evidence we actually have.

Utilities and expectations

- A **utility** is a number attached to an outcome, saying how good it is by your lights.

Utilities and expectations

- A **utility** is a number attached to an outcome, saying how good it is by your lights.
- An **expectation** of a number is just its **average** value.

Utilities and expectations

- A **utility** is a number attached to an outcome, saying how good it is by your lights.
- An **expectation** of a number is just its **average** value.

Example

\$100 if a fair coin lands heads, \$0 if tails. The **expected** amount is **\$50**.

Utilities and expectations

- A **utility** is a number attached to an outcome, saying how good it is by your lights.
- An **expectation** of a number is just its **average** value.

Example

$100 if a fair coin lands heads, $0 if tails. The **expected** amount is **$50**.

Note: the expected amount is $50 even though **I know for certain I won't get $50**.

Maximising expected utility

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

Maximising expected utility

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

- Each row is a set of possible outcomes with probabilities attached.

Maximising expected utility

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

- Each row is a set of possible outcomes with probabilities attached.
- So each **act** has an **expected utility**.

It depends how you assign the utilities

Risk-neutral: every dollar worth the same

$u(\$100) = 100, \quad u(\$50) = 50$

It depends how you assign the utilities

Risk-neutral: every dollar worth the same

$u(\$100) = 100, \quad u(\$50) = 50$

Bet: **50**. Decline: **50**. **Indifferent.**

It depends how you assign the utilities

Risk-neutral: every dollar worth the same

$u(\$100) = 100, \quad u(\$50) = 50$

Bet: **50**. Decline: **50**. **Indifferent.**

Risk-averse: \$100 isn't worth twice \$50

$u(\$100) = 80, \quad u(\$50) = 50$

It depends how you assign the utilities

Risk-neutral: every dollar worth the same

$u(\$100) = 100, \quad u(\$50) = 50$

Bet: **50**. Decline: **50**. **Indifferent.**

Risk-averse: \$100 isn't worth twice \$50

$u(\$100) = 80, \quad u(\$50) = 50$

Bet: **40**. Decline: **50**. **Don't bet.**

It depends how you assign the utilities

Risk-neutral: every dollar worth the same

$u(\$100) = 100, \quad u(\$50) = 50$

Bet: **50**. Decline: **50**. **Indifferent.**

Risk-averse: \$100 isn't worth twice \$50

$u(\$100) = 80, \quad u(\$50) = 50$

Bet: **40**. Decline: **50**. **Don't bet.**

Both are ways of maximising expected utility.

One thing I haven't done

I haven't given you **any argument** that you *should* maximise expected utility.

One thing I haven't done

I haven't given you **any argument** that you *should* maximise expected utility.

That comes in several weeks' time.

Which credences should we have?

We won't answer that fully. But there's a partial answer: **Bayesianism**.

Which credences should we have?

We won't answer that fully. But there's a partial answer: **Bayesianism**.

Probabilism

Credences must obey the **probability axioms**.

Which credences should we have?

We won't answer that fully. But there's a partial answer: **Bayesianism**.

Probabilism

Credences must obey the **probability axioms**.

The Principle of Conditionalisation

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

The probability axioms

- Write X for the set of all possible outcomes — the event that *something happens*.

The probability axioms

- Write X for the set of all possible outcomes — the event that *something happens*.
- An **event** is a set of possible outcomes. (Including none of them, and all of them.)

The probability axioms

- Write X for the set of all possible outcomes — the event that *something happens*.
- An **event** is a set of possible outcomes. (Including none of them, and all of them.)
- Write $\mathrm{Cr}(\sigma)$ for your credence in event σ.

The probability axioms

- Write X for the set of all possible outcomes — the event that *something happens*.
- An **event** is a set of possible outcomes. (Including none of them, and all of them.)
- Write $\mathrm{Cr}(\sigma)$ for your credence in event σ.

The axioms

(i) $\mathrm{Cr}(\sigma) \geq 0$

The probability axioms

- Write X for the set of all possible outcomes — the event that *something happens*.
- An **event** is a set of possible outcomes. (Including none of them, and all of them.)
- Write $\mathrm{Cr}(\sigma)$ for your credence in event σ.

The axioms

(i) $\mathrm{Cr}(\sigma) \geq 0$

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

The probability axioms

- Write X for the set of all possible outcomes — the event that *something happens*.
- An **event** is a set of possible outcomes. (Including none of them, and all of them.)
- Write $\mathrm{Cr}(\sigma)$ for your credence in event σ.

The axioms

(i) $\mathrm{Cr}(\sigma) \geq 0$

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

(iii) If σ and τ are **disjoint**, then $\mathrm{Cr}(\sigma \cup \tau) = \mathrm{Cr}(\sigma) + \mathrm{Cr}(\tau)$

The probability axioms

- Write X for the set of all possible outcomes — the event that *something happens*.
- An **event** is a set of possible outcomes. (Including none of them, and all of them.)
- Write $\mathrm{Cr}(\sigma)$ for your credence in event σ.

The axioms

(i) $\mathrm{Cr}(\sigma) \geq 0$

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

(iii) If σ and τ are **disjoint**, then $\mathrm{Cr}(\sigma \cup \tau) = \mathrm{Cr}(\sigma) + \mathrm{Cr}(\tau)$

Question: Can you show that $\mathrm{Cr}(\sigma) + \mathrm{Cr}(\sigma^c) = 1$?

Left for next week: Conditionalisation

A sneak peek:

Left for next week: Conditionalisation

A sneak peek:

Conditional probability formula

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

For next week

Reading

BDRC, chapters 2 and 3.

For next week

Reading

BDRC, chapters 2 and 3.

- Chapter 1 is also worth a look, to brush up on today.

For next week

Reading

BDRC, chapters 2 and 3.

- Chapter 1 is also worth a look, to brush up on today.
- **Register for a tutorial** if you haven't — they start in week 4.

For next week

Reading

BDRC, chapters 2 and 3.

- Chapter 1 is also worth a look, to brush up on today.
- **Register for a tutorial** if you haven't — they start in week 4.

See you next week.