Introduction to Decision Theory

Conditionals in Decision Theory
Lecture 1

Calum McNamara and Paolo Santorio
Indiana University, Bloomington, and University of Maryland, College Park

ESSLLI 2026, Prague

Introduction: conditionals in deliberation

A decision problem

Suppose we offer you the following bet:

	Heads	*Tails*
Accept	\$5	\$5
Decline	\$0	\$0

What should you do, and why?

A decision problem

If you have responded the way we wanted you too, you will have said something like:

A decision problem

If you have responded the way we wanted you too, you will have said something like:

(i) If I accept, I will win \$5 no matter what.

A decision problem

If you have responded the way we wanted you too, you will have said something like:

(i) If I accept, I will win $5 no matter what.
(ii) If I decline, I will win nothing.

A decision problem

If you have responded the way we wanted you too, you will have said something like:

(i) If I accept, I will win \$5 no matter what.

(ii) If I decline, I will win nothing.

(i) and (ii) are conditionals of deliberation.

A decision problem

If you have responded the way we wanted you too, you will have said something like:

(i) If I accept, I will win $5 no matter what.

(ii) If I decline, I will win nothing.

(i) and (ii) are conditionals of deliberation.

- They are used to specify connections between possible actions that you can take and possible outcomes.

A decision problem

If you have responded the way we wanted you too, you will have said something like:

(i) If I accept, I will win $5 no matter what.
(ii) If I decline, I will win nothing.

(i) and (ii) are conditionals of deliberation.

- They are used to specify connections between possible actions that you can take and possible outcomes.
- Arguably, they are central to the way we reason about what to do.

A decision problem

If you have responded the way we wanted you too, you will have said something like:

(i) If I accept, I will win $5 no matter what.

(ii) If I decline, I will win nothing.

(i) and (ii) are conditionals of deliberation.

- They are used to specify connections between possible actions that you can take and possible outcomes.
- Arguably, they are central to the way we reason about what to do.

This class is about these conditionals, and how they interact with the standard formal framework for rational decision making.

Conditionals in deliberation

By a conditional, we mean an 'If... then... sentence' (or a proposition expressed by such a sentence).

[1] If you don't know this example already, you will very soon.

Conditionals in deliberation

By a conditional, we mean an ‘If… then… sentence’ (or a proposition expressed by such a sentence).

Clearly, when we deliberate about what to do, we entertain such sentences—we just saw two examples. Here are some more:

[1] If you don’t know this example already, you will very soon.

Conditionals in deliberation

By a conditional, we mean an 'If... then... sentence' (or a proposition expressed by such a sentence).

Clearly, when we deliberate about what to do, we entertain such sentences—we just saw two examples. Here are some more:

(1) a. Should I take my umbrella? If I take it, I'll stay dry; but I'll also have to carry it around all day.

[1] If you don't know this example already, you will very soon.

Conditionals in deliberation

By a conditional, we mean an 'If... then... sentence' (or a proposition expressed by such a sentence).

Clearly, when we deliberate about what to do, we entertain such sentences—we just saw two examples. Here are some more:

(1) a. Should I take my umbrella? If I take it, I'll stay dry; but I'll also have to carry it around all day.

b. If I cross the river here, I'll get wet; but if I head further upstream, I'll be late (Williamson, 2020).

[1] If you don't know this example already, you will very soon.

Conditionals in deliberation

By a conditional, we mean an 'If… then… sentence' (or a proposition expressed by such a sentence).

Clearly, when we deliberate about what to do, we entertain such sentences—we just saw two examples. Here are some more:

(1) a. Should I take my umbrella? If I take it, I'll stay dry; but I'll also have to carry it around all day.

b. If I cross the river here, I'll get wet; but if I head further upstream, I'll be late (Williamson, 2020).

c. If I take just the opaque box, I'll be rich; but if I take both boxes, I'll be poor.[1]

[1] If you don't know this example already, you will very soon.

English language conditionals come in two broad families:

Indicatives and subjunctives

English language conditionals come in two broad families:

Indicative	If A is true, C is true.
Subjunctive/counterfactual	If A were true, C would be true.

A minimal pair

To illustrate the difference, consider a famous minimal pair, from Bennett (1988):

A minimal pair

To illustrate the difference, consider a famous minimal pair, from Bennett (1988):

Kind	Sentence	Natural verdict
Indicative	If Shakespeare didn't write *Hamlet*, someone else did.	Very likely true
Subjunctive	If Shakespeare hadn't written *Hamlet*, someone else would have.	Very likely false

A minimal pair

To illustrate the difference, consider a famous minimal pair, from Bennett (1988):

Kind	**Sentence**	**Natural verdict**
Indicative	If Shakespeare didn't write *Hamlet*, someone else did.	Very likely true
Subjunctive	If Shakespeare hadn't written *Hamlet*, someone else would have.	Very likely false

Here, we have the same antecedent, but different conditional mood, and different truth conditions (intuitively).

A rough working contrast

The differences between indicatives and subjunctives seem to be as follows:

Indicatives

- Epistemic or matter-of-fact.
- Hold fixed what is known or taken for granted.

A rough working contrast

The differences between indicatives and subjunctives seem to be as follows:

Indicatives

- Epistemic or matter-of-fact.
- Hold fixed what is known or taken for granted.

Subjunctives

- Causal, nomological, or metaphysical.
- Hold fixed objective dependencies.

A rough working contrast

The differences between indicatives and subjunctives seem to be as follows:

Indicatives

- Epistemic or matter-of-fact.
- Hold fixed what is known or taken for granted.

Subjunctives

- Causal, nomological, or metaphysical.
- Hold fixed objective dependencies.

(This is only a first pass, but it is enough for present purposes.)

A puzzle

The fact, however, that conditionals come in two varieties raises a puzzle, when we think about their use in deliberation.

A puzzle

The fact, however, that conditionals come in two varieties raises a puzzle, when we think about their use in deliberation.

Question: Are the "conditionals of deliberation" indicatives? Subjunctives? Both?

A puzzle

The fact, however, that conditionals come in two varieties raises a puzzle, when we think about their use in deliberation.

Question: Are the "conditionals of deliberation" indicatives? Subjunctives? Both?

It turns out that authors are divided here:

- Some, like Gibbard and Harper (1978), say: the conditionals of deliberation are subjunctives.

A puzzle

The fact, however, that conditionals come in two varieties raises a puzzle, when we think about their use in deliberation.

Question: Are the "conditionals of deliberation" indicatives? Subjunctives? Both?

It turns out that authors are divided here:

- Some, like Gibbard and Harper (1978), say: the conditionals of deliberation are subjunctives.
- Others, like Jeffrey (1983) and DeRose (2010), say: the conditionals of deliberation are indicatives.

Decision theory

This debate makes its way into decision theory.

[2] Again, if you don't know *Newcomb's problem*, then you will very soon.

Decision theory

This debate makes its way into decision theory.

For example, Lewis (1981, p. 377) characterized the debate between *causal* and *evidential decision theorists* as a debate about which conditionals are relevant to decision-making:[2]

[2] Again, if you don't know *Newcomb's problem*, then you will very soon.

Decision theory

This debate makes its way into decision theory.

For example, Lewis (1981, p. 377) characterized the debate between *causal* and *evidential decision theorists* as a debate about which conditionals are relevant to decision-making:[2]

> *Some think that in (a suitable version of) Newcomb's problem, it is rational to take only one box. These one-boxers think of the situation as a choice between a million and a thousand. They [viz., evidential decision theorists]* **are convinced by indicative conditionals: if I take one box I will be a millionaire, but if I take both boxes I will not***. Others, and I for one, think it rational to take both boxes. We two-boxers think that whether the million already awaits us or not, we have no choice between taking it and leaving it. We [causal decision theorists]* **are convinced by counterfactual conditionals: If I took only one box, I would be poorer by a thousand than I will be after taking both.**

[2] Again, if you don't know *Newcomb's problem*, then you will very soon.

Matters arising

Let us call a conditional a conditional of deliberation when:

1. it is future-directed; and
2. its antecedent concerns an option now available to you.

Matters arising

Let us call a conditional a conditional of deliberation when:

1. it is future-directed; and
2. its antecedent concerns an option now available to you.

Indicative	If I take my umbrella, I will stay dry.
Subjunctive	If I were to take my umbrella, I would stay dry.

Matters arising

Let us call a conditional a conditional of deliberation when:

1. it is future-directed; and
2. its antecedent concerns an option now available to you.

Indicative	If I take my umbrella, I will stay dry.
Subjunctive	If I were to take my umbrella, I would stay dry.

(Notice that the contrast here looks much less sharp than in Bennett's Shakespeare example.)

So: which is it? Are the conditionals of deliberation indicatives, or subjunctives?

We have a view about this—our view will come up in due course. However, it turns out that, even once we've answered this question, issues remain:

Semantics

We have a view about this—our view will come up in due course. However, it turns out that, even once we've answered this question, issues remain:

For example, we still need to say what are the semantics for the conditionals of deliberation:

- This is pressing because—as we will see—different semantics lead to better and worse recommendations in concrete decision problems.

Semantics

We have a view about this—our view will come up in due course. However, it turns out that, even once we've answered this question, issues remain:

For example, we still need to say what are the semantics for the conditionals of deliberation:

- This is pressing because—as we will see—different semantics lead to better and worse recommendations in concrete decision problems.
- Indeed, in Lecture 3 we will see that the standard semantics for counterfactuals leads us to deliberate as if we could break the laws of nature.

Probability

Additionally, it's fairly clear that, if conditionals are going to feature in deliberation, then we need to attach probabilities to them.

- For example, if you're highly confident in 'If I take this bet, I'll win', it seems rational to take the bet; less so if you're not highly confident.

Probability

Additionally, it's fairly clear that, if conditionals are going to feature in deliberation, then we need to attach probabilities to them.

- For example, if you're highly confident in 'If I take this bet, I'll win', it seems rational to take the bet; less so if you're not highly confident.

However, the probabilities of conditionals is an area that's fraught with difficulty:

- Consider: 'If that (fair, six-sided) die landed on an even number, then it landed on 2'.

Probability

Additionally, it's fairly clear that, if conditionals are going to feature in deliberation, then we need to attach probabilities to them.

- For example, if you're highly confident in 'If I take this bet, I'll win', it seems rational to take the bet; less so if you're not highly confident.

However, the probabilities of conditionals is an area that's fraught with difficulty:

- Consider: 'If that (fair, six-sided) die landed on an even number, then it landed on 2'.
- How confident should you be that this conditional is true?

Probability

Additionally, it's fairly clear that, if conditionals are going to feature in deliberation, then we need to attach probabilities to them.

- For example, if you're highly confident in 'If I take this bet, I'll win', it seems rational to take the bet; less so if you're not highly confident.

However, the probabilities of conditionals is an area that's fraught with difficulty:

- Consider: 'If that (fair, six-sided) die landed on an even number, then it landed on 2'.
- How confident should you be that this conditional is true?
- Intuition says '1/3' since that is the probability that the die landed on 2, *given* that it landed on even—a conditional probability.

Probability

Additionally, it's fairly clear that, if conditionals are going to feature in deliberation, then we need to attach probabilities to them.

- For example, if you're highly confident in 'If I take this bet, I'll win', it seems rational to take the bet; less so if you're not highly confident.

However, the probabilities of conditionals is an area that's fraught with difficulty:

- Consider: 'If that (fair, six-sided) die landed on an even number, then it landed on 2'.
- How confident should you be that this conditional is true?
- Intuition says '1/3' since that is the probability that the die landed on 2, *given* that it landed on even—a conditional probability.
- But famous results from Lewis (1976) show that a very general answer along these lines is inconsistent with standard assumptions about probability!

Lecture 1: decision theory

On the agenda today

Today will be a broad introduction to decision theory. We will look at:

On the agenda today

Today will be a broad introduction to decision theory. We will look at:

1. Expected utility theory.

On the agenda today

Today will be a broad introduction to decision theory. We will look at:

1. Expected utility theory.
2. The *act-state dependence problem*.

On the agenda today

Today will be a broad introduction to decision theory. We will look at:

1. Expected utility theory.
2. The *act-state dependence problem*.
3. Evidential decision theory.

On the agenda today

Today will be a broad introduction to decision theory. We will look at:

1. Expected utility theory.
2. The *act-state dependence problem*.
3. Evidential decision theory.
4. Newcomb's problem.

On the agenda today

Today will be a broad introduction to decision theory. We will look at:

1. Expected utility theory.
2. The *act-state dependence problem*.
3. Evidential decision theory.
4. Newcomb's problem.
5. Causal decision theory.

On the agenda today

Today will be a broad introduction to decision theory. We will look at:

1. Expected utility theory.
2. The *act-state dependence problem*.
3. Evidential decision theory.
4. Newcomb's problem.
5. Causal decision theory.
6. Suppositional decision theories.

On the agenda today

Today will be a broad introduction to decision theory. We will look at:

1. Expected utility theory.
2. The *act-state dependence problem*.
3. Evidential decision theory.
4. Newcomb's problem.
5. Causal decision theory.
6. Suppositional decision theories.

For many of you, some of what we do today will be review. Stick with us.

Formal preliminaries

We need to start by introducing some formalism:

- $\mathcal{W} = \{w_1, \ldots, w_n\}$ is a finite set of possible worlds.[3]

[3] Our assumption that $\mathcal{W}$ is finite is made just to keep the mathematics simpler.

Formal preliminaries

We need to start by introducing some formalism:

- $\mathcal{W} = \{w_1, \ldots, w_n\}$ is a finite set of possible worlds.[3]
- Propositions are subsets of $\mathcal{W}$; so the set of all propositions is $\mathcal{P}(\mathcal{W})$ (the power set of $\mathcal{W}$).

[3] Our assumption that $\mathcal{W}$ is finite is made just to keep the mathematics simpler.

Formal preliminaries

We need to start by introducing some formalism:

- $\mathcal{W} = \{w_1, \ldots, w_n\}$ is a finite set of possible worlds.[3]
- Propositions are subsets of $\mathcal{W}$; so the set of all propositions is $\mathcal{P}(\mathcal{W})$ (the power set of $\mathcal{W}$).
- $\mathcal{O}$ is a partition of options: propositions you can *make* true by deciding.

[3] Our assumption that $\mathcal{W}$ is finite is made just to keep the mathematics simpler.

Formal preliminaries

We need to start by introducing some formalism:

- $\mathcal{W} = \{w_1, \ldots, w_n\}$ is a finite set of possible worlds.[3]
- Propositions are subsets of $\mathcal{W}$; so the set of all propositions is $\mathcal{P}(\mathcal{W})$ (the power set of $\mathcal{W}$).
- $\mathcal{O}$ is a partition of options: propositions you can *make* true by deciding.
- $\mathcal{S}$ is a partition of state propositions.

[3] Our assumption that $\mathcal{W}$ is finite is made just to keep the mathematics simpler.

Formal preliminaries

We need to start by introducing some formalism:

- $\mathcal{W} = \{w_1, \ldots, w_n\}$ is a finite set of possible worlds.[3]
- Propositions are subsets of $\mathcal{W}$; so the set of all propositions is $\mathcal{P}(\mathcal{W})$ (the power set of $\mathcal{W}$).
- $\mathcal{O}$ is a partition of options: propositions you can *make* true by deciding.
- $\mathcal{S}$ is a partition of state propositions.
- Outcomes are intersections $O \cap S$ of an option and a state.

[3]Our assumption that $\mathcal{W}$ is finite is made just to keep the mathematics simpler.

Formal preliminaries

We need to start by introducing some formalism:

- $\mathcal{W} = \{w_1, \ldots, w_n\}$ is a finite set of possible worlds.[3]
- Propositions are subsets of $\mathcal{W}$; so the set of all propositions is $\mathcal{P}(\mathcal{W})$ (the power set of $\mathcal{W}$).
- $\mathcal{O}$ is a partition of options: propositions you can *make* true by deciding.
- $\mathcal{S}$ is a partition of state propositions.
- Outcomes are intersections $O \cap S$ of an option and a state.
- p is your probability function, and u is your utility function.

[3] Our assumption that $\mathcal{W}$ is finite is made just to keep the mathematics simpler.

Formal preliminaries

We need to start by introducing some formalism:

- $\mathcal{W} = \{w_1, \ldots, w_n\}$ is a finite set of possible worlds.[3]
- Propositions are subsets of $\mathcal{W}$; so the set of all propositions is $\mathcal{P}(\mathcal{W})$ (the power set of $\mathcal{W}$).
- $\mathcal{O}$ is a partition of options: propositions you can *make* true by deciding.
- $\mathcal{S}$ is a partition of state propositions.
- Outcomes are intersections $O \cap S$ of an option and a state.
- p is your probability function, and u is your utility function.

(Thus, we assume that states are fine-grained enough that each option-state intersection settles everything you care about, at least at the time in question.)

[3] Our assumption that $\mathcal{W}$ is finite is made just to keep the mathematics simpler.

Expected utility theory

Your utility function assigns a value (real number) to *outcomes*. We can also assign a value to *options*. For now, we take the correct such value to be expected utility (EU):

Your utility function assigns a value (real number) to *outcomes*. We can also assign a value to *options*. For now, we take the correct such value to be expected utility (EU):

Expected utility

$$\mathrm{EU}(O) = \sum_i p(S_i) \cdot u(O \cap S_i)$$

Expected utility theory

Your utility function assigns a value (real number) to *outcomes*. We can also assign a value to *options*. For now, we take the correct such value to be expected utility (EU):

Expected utility

$$\mathrm{EU}(O) = \sum_i p(S_i) \cdot u(O \cap S_i)$$

EU theory then says: when you're choosing between the option-propositions, O, you should choose one that maximizes EU.

Example

EU-maximization looks good in many cases. For instance:

	Heads	*Tails*
Accept	\$2	−\$1
Decline	\$0	\$0

Example

EU-maximization looks good in many cases. For instance:

	Heads	*Tails*
Accept	\$2	−\$1
Decline	\$0	\$0

The EU calculations:

$$\text{EU}(\textit{Accept}) = (1/2)(2) + (1/2)(-1) = 1/2$$
$$\text{EU}(\textit{Decline}) = (1/2)(0) + (1/2)(0) = 0$$

Example

EU-maximization looks good in many cases. For instance:

	Heads	*Tails*
Accept	\$2	−\$1
Decline	\$0	\$0

The EU calculations:

$$\text{EU}(\textit{Accept}) = (1/2)(2) + (1/2)(-1) = 1/2$$
$$\text{EU}(\textit{Decline}) = (1/2)(0) + (1/2)(0) = 0$$

So, $\text{EU}(\textit{Accept}) > \text{EU}(\textit{Decline})$, and EU theory says you should accept.

The Extortionist

To this day, EU theory is the standard theory of rational decision-making, in economics and other fields.

The Extortionist

To this day, EU theory is the standard theory of rational decision-making, in economics and other fields.

But as stated, it can't be completely correct.

Consider another example (Joyce, 1999):

The Extortionist

To this day, EU theory is the standard theory of rational decision-making, in economics and other fields.

But as stated, it can't be completely correct.

Consider another example (Joyce, 1999):

> *Suppose you have just parked in a seedy neighborhood when a man approaches and offers to "protect" your car from harm for \$10. You recognize this as extortion and have heard that people who refuse "protection" invariably return to find their windshields smashed. Those who pay find their cars intact. You cannot park anywhere else because you are late for an important meeting. It costs \$400 to replace a windshield. Should you buy "protection"? (pp. 115–16)*

The decision table

Here is the table for Joyce's decision problem:

	Smashed	*Not smashed*
Pay	−$410	−$10
Don't pay	−$400	$0

The decision table

Here is the table for Joyce's decision problem:

	Smashed	*Not smashed*
Pay	$-$410	$-$10
Don't pay	$-$400	$0

Exercise! What does EU theory say you should do in this situation?

(*Pssst!* Do you even need to do any calculations?)

The problem

Answer: EU theory says that you should decline to pay the extortionist.

The problem

Answer: EU theory says that you should decline to pay the extortionist.

But that seems wrong. As Joyce himself describes the issue:

> *[In The Extortionist] Your choice has a* **direct influence** *on the state of the world; refusing to pay* **makes it likely** *that your windshield will be smashed while paying* **makes this unlikely**. *The extortionist is a despicable person, but he has you over a barrel and investing a mere* \$10 *now saves* \$400 *down the line. You should pay now (and alert the police later). (p. 116, with emphasis added)*

The problem

Answer: EU theory says that you should decline to pay the extortionist.

But that seems wrong. As Joyce himself describes the issue:

> *[In The Extortionist] Your choice has a* **direct influence** *on the state of the world; refusing to pay* **makes it likely** *that your windshield will be smashed while paying* **makes this unlikely**. *The extortionist is a despicable person, but he has you over a barrel and investing a mere* \$10 *now saves* \$400 *down the line. You should pay now (and alert the police later). (p. 116, with emphasis added)*

Thus, the issue with standard EU theory is that it doesn't take into account how states can depend upon your choices.

The result is that, sometimes, the theory gives the wrong results.

Evidential decision theory as a solution

What we need is a decision theory which drops this implausible assumption of "act-state independence".

Evidential decision theory as a solution

What we need is a decision theory which drops this implausible assumption of "act-state independence".

The first theory proposed in this vein was formulated by Richard Jeffrey (1983).

Evidential decision theory as a solution

What we need is a decision theory which drops this implausible assumption of "act-state independence".

The first theory proposed in this vein was formulated by Richard Jeffrey (1983).

The theory is now known as evidential decision theory (EDT).

Jeffrey's idea

Jeffrey's key idea was to replace unconditional probabilities in the definition of EU with conditional probabilities:[4]

[4] We define conditional probabilities using the *ratio formula:* $p(C \mid A) := p(A \cap C)/p(A)$, *assuming* $p(A) > 0$.

Jeffrey's idea

Jeffrey's key idea was to replace unconditional probabilities in the definition of EU with conditional probabilities:[4]

Evidential expected utility

$$\mathrm{EEU}(O) = \sum_i p(S_i \mid O) \cdot u(O \cap S_i)$$

[4] We define conditional probabilities using the *ratio formula:* $p(C \mid A) := p(A \cap C)/p(A)$, *assuming* $p(A) > 0$.

Jeffrey's idea

Jeffrey's key idea was to replace unconditional probabilities in the definition of EU with conditional probabilities:[4]

Evidential expected utility

$$\mathrm{EEU}(O) = \sum_i p(S_i \mid O) \cdot u(O \cap S_i)$$

The idea is that $p(S \mid O)$ measures how much evidence you think choosing O would provide that state S obtains.

[4] We define conditional probabilities using the *ratio formula:* $p(C \mid A) := p(A \cap C)/p(A)$, *assuming* $p(A) > 0$.

Jeffrey's idea

Jeffrey's key idea was to replace unconditional probabilities in the definition of EU with conditional probabilities:[4]

Evidential expected utility

$$\mathrm{EEU}(O) = \sum_i p(S_i \mid O) \cdot u(O \cap S_i)$$

The idea is that $p(S \mid O)$ measures how much evidence you think choosing O would provide that state S obtains.

The resulting theory—evidential decision theory—says: choose an option that has the highest evidential expected utility (EEU).

[4] We define conditional probabilities using the *ratio formula:* $p(C \mid A) := p(A \cap C)/p(A)$, *assuming* $p(A) > 0$.

EDT and The Extortionist

Jeffrey's theory gets the right result in cases like Joyce's *Extortionist*.

EDT and The Extortionist

Jeffrey's theory gets the right result in cases like Joyce's *Extortionist*.

To see this, assume:

$$p(\textit{Not smashed} \mid \textit{Pay}) \approx 1$$
$$p(\textit{Smashed} \mid \textit{Don't pay}) \approx 1$$

EDT and The Extortionist

Jeffrey's theory gets the right result in cases like Joyce's *Extortionist*.

To see this, assume:

$$p(\textit{Not smashed} \mid \textit{Pay}) \approx 1$$
$$p(\textit{Smashed} \mid \textit{Don't pay}) \approx 1$$

Then:

$$\text{EEU}(\textit{Pay}) \approx (0)(-410) + (1)(-10) = -10$$
$$\text{EEU}(\textit{Don't pay}) \approx (1)(-400) + (0)(0) = -400$$

EDT and The Extortionist

Jeffrey's theory gets the right result in cases like Joyce's *Extortionist*.

To see this, assume:

$$p(\textit{Not smashed} \mid \textit{Pay}) \approx 1$$
$$p(\textit{Smashed} \mid \textit{Don't pay}) \approx 1$$

Then:

$$\text{EEU}(\textit{Pay}) \approx (0)(-410) + (1)(-10) = -10$$
$$\text{EEU}(\textit{Don't pay}) \approx (1)(-400) + (0)(0) = -400$$

Thus, EDT recommends paying, because paying is excellent evidence that the windshield will remain intact. This, in turn, is the right answer.

Digression: world-based notation

A brief digression: earlier, we defined your utility function as mapping *outcomes* to real numbers.

Digression: world-based notation

A brief digression: earlier, we defined your utility function as mapping *outcomes* to real numbers.

For EDT (and all the theories that follow from now on) we can equivalently let utilities attach directly to worlds:

$$u : \mathcal{W} \to \mathbb{R}$$

Digression: world-based notation

A brief digression: earlier, we defined your utility function as mapping *outcomes* to real numbers.

For EDT (and all the theories that follow from now on) we can equivalently let utilities attach directly to worlds:

$$u : \mathcal{W} \to \mathbb{R}$$

Then:

Evidential expected utility (worlds version)

$$\mathrm{EEU}(O) = \sum_{w} p(w \mid O) \cdot u(w)$$

Digression: world-based notation

A brief digression: earlier, we defined your utility function as mapping *outcomes* to real numbers.

For EDT (and all the theories that follow from now on) we can equivalently let utilities attach directly to worlds:

$$u : \mathcal{W} \to \mathbb{R}$$

Then:

Evidential expected utility (worlds version)

$$\mathrm{EEU}(O) = \sum_{w} p(w \mid O) \cdot u(w)$$

This formulation of EEU is sometimes easier to use—we will switch back and forth liberally.

(As an exercise, however, it's worth asking why we can't define u like this, in the case of EU theory.)

Newcomb's problem

Thus, EDT looks like it's an improvement on ordinary EU theory—and it is.

[5] This is the decision problem we hinted about at the outset. It was first introduced into the literature by Nozick (1969), but attributed to the physicist William Newcomb.

Newcomb's problem

Thus, EDT looks like it's an improvement on ordinary EU theory—and it is.

But there are some cases in which it, too, seems to get the wrong results—at least according to many authors:[5]

[5] This is the decision problem we hinted about at the outset. It was first introduced into the literature by Nozick (1969), but attributed to the physicist William Newcomb.

Thus, EDT looks like it's an improvement on ordinary EU theory—and it is.

But there are some cases in which it, too, seems to get the wrong results—at least according to many authors:[5]

> *Newcomb. In front of you are two boxes, labeled 'A' and 'B', respectively. You can take either just box A, or both boxes. Box B is transparent and contains what you can see to be a* $1,000 *bill. Box A, in contrast, is opaque, and you don't know what's inside it. The contents of box A were determined yesterday on the basis of a prediction I made about your behavior. If I predicted that you'd take only box A, then I put* $1,000,000 *inside that box. However, if I predicted that you'd take both boxes, then I left the opaque box empty. Note that I'm a highly reliable predictor of your behavior. So, with that in mind: What is your choice?*

[5] This is the decision problem we hinted about at the outset. It was first introduced into the literature by Nozick (1969), but attributed to the physicist William Newcomb.

EDT in Newcomb

Here is the corresponding decision table:

	Million	*No million*
One box	\$1,000,000	\$0
Both boxes	\$1,001,000	\$1,000

EDT in Newcomb

Here is the corresponding decision table:

	Million	*No million*
One box	\$1,000,000	\$0
Both boxes	\$1,001,000	\$1,000

What does EDT recommend? Here, the calculations are more-or-less identical to those in *The Extortionist*.

EDT in Newcomb

Here is the corresponding decision table:

	Million	*No million*
One box	\$1,000,000	\$0
Both boxes	\$1,001,000	\$1,000

What does EDT recommend? Here, the calculations are more-or-less identical to those in *The Extortionist*.

$$p(\textit{Million} \mid \textit{One box}) \approx 1$$
$$p(\textit{No million} \mid \textit{Both boxes}) \approx 1$$

EDT in Newcomb

Here is the corresponding decision table:

	Million	*No million*
One box	\$1,000,000	\$0
Both boxes	\$1,001,000	\$1,000

What does EDT recommend? Here, the calculations are more-or-less identical to those in *The Extortionist*.

$$p(\textit{Million} \mid \textit{One box}) \approx 1$$
$$p(\textit{No million} \mid \textit{Both boxes}) \approx 1$$

So:

$$\text{EEU}(\textit{One box}) \approx (1)(1,000,000) + (0)(0) = 1,000,000$$
$$\text{EEU}(\textit{Both boxes}) \approx (0)(1,001,000) + (1)(1,000) = 1,000$$

Is one-boxing right?

Thus, EDT recommends "one-boxing".

Is one-boxing right?

Thus, EDT recommends "one-boxing".

No doubt, some of you will think this is correct. If so, you are "convinced by indicative conditionals", as Lewis earlier remarked:

Is one-boxing right?

Thus, EDT recommends "one-boxing".

No doubt, some of you will think this is correct. If so, you are "convinced by indicative conditionals", as Lewis earlier remarked:

> *Some think that in... Newcomb's problem, it is rational to take only one box. These one-boxers think of the situation as a choice between a million and a thousand. They are convinced by indicative conditionals: if I take one box I will be a millionaire, but if I take both boxes I will not.*

Is one-boxing right?

Thus, EDT recommends "one-boxing".

No doubt, some of you will think this is correct. If so, you are "convinced by indicative conditionals", as Lewis earlier remarked:

> *Some think that in... Newcomb's problem, it is rational to take only one box. These one-boxers think of the situation as a choice between a million and a thousand. They are convinced by indicative conditionals: if I take one box I will be a millionaire, but if I take both boxes I will not.*

(*Pssst!* Note that Lewis is talking here about indicative conditionals; but Jeffrey's theory uses *conditional probabilities*. Thus, Lewis seems to have in mind the thesis that $p(S \mid O) = p(O \to S)$, where $O \to S$ is the *indicative conditional*. More on this relationship in Lecture 4.)

Is two-boxing right?

Many philosophers have argued, however—including Lewis, and me—that EDT gets things wrong here. One-boxing is not the rational choice.

Why? Because...

Is two-boxing right?

Many philosophers have argued, however—including Lewis, and me—that EDT gets things wrong here. One-boxing is not the rational choice.

Why? Because...

- ...the money was placed in the opaque box yesterday.

Is two-boxing right?

Many philosophers have argued, however—including Lewis, and me—that EDT gets things wrong here. One-boxing is not the rational choice.

Why? Because...

- ...the money was placed in the opaque box yesterday.
- ...there's nothing you do now that can *causally affect* whether it is there or not.

Is two-boxing right?

Many philosophers have argued, however—including Lewis, and me—that EDT gets things wrong here. One-boxing is not the rational choice.

Why? Because...

- ...the money was placed in the opaque box yesterday.
- ...there's nothing you do now that can *causally affect* whether it is there or not.
- ...taking both boxes gets you $1,000 more than one-boxing would, *no matter what the predictor predicted*.

Is two-boxing right?

Many philosophers have argued, however—including Lewis, and me—that EDT gets things wrong here. One-boxing is not the rational choice.

Why? Because...

- ...the money was placed in the opaque box yesterday.
- ...there's nothing you do now that can *causally affect* whether it is there or not.
- ...taking both boxes gets you $1,000 more than one-boxing would, *no matter what the predictor predicted*.

Diagnosis: EDT tracks evidential dependence; but rational decision-making should track causal dependence.

Causal decision theory

Causal decision theory was introduced as an alternative to EDT, in order to adequately handle cases like Newcomb.

Causal decision theory

Causal decision theory was introduced as an alternative to EDT, in order to adequately handle cases like Newcomb.

There are different versions of this theory on the market; but for this course, we will focus on a version introduced by Stalnaker (1981) and developed by Gibbard and Harper (1978):

Causal decision theory

Causal decision theory was introduced as an alternative to EDT, in order to adequately handle cases like Newcomb.

There are different versions of this theory on the market; but for this course, we will focus on a version introduced by Stalnaker (1981) and developed by Gibbard and Harper (1978):

This theory appeals, not to conditional probabilities (as in EDT), but to probabilities of counterfactual conditionals.

Stalnaker-style CDT

Let $O \mathrel{\Box\!\!\rightarrow} S$ symbolize the proposition expressed by: 'If I were to choose O, then state S would obtain'.

[6] Exercise: give the "states" formulation of causal expected utility (CEU), instead of the "worlds" formulation.

Stalnaker-style CDT

Let $O \,\Box\!\!\rightarrow S$ symbolize the proposition expressed by: 'If I were to choose O, then state S would obtain'.

We call $\Box\!\!\rightarrow$ the counterfactual (or subjunctive) conditional operator.

[6] Exercise: give the "states" formulation of causal expected utility (CEU), instead of the "worlds" formulation.

Stalnaker-style CDT

Let $O \mathrel{\Box\!\!\rightarrow} S$ symbolize the proposition expressed by: 'If I were to choose O, then state S would obtain'.

We call $\mathrel{\Box\!\!\rightarrow}$ the counterfactual (or subjunctive) conditional operator.

Then:[6]

Causal expected utility

$$\mathrm{CEU}(O) = \sum_{w} p(O \mathrel{\Box\!\!\rightarrow} w) \cdot u(w)$$

[6] Exercise: give the "states" formulation of causal expected utility (CEU), instead of the "worlds" formulation.

Stalnaker-style CDT

Let $O \mathrel{\Box\!\!\rightarrow} S$ symbolize the proposition expressed by: 'If I were to choose O, then state S would obtain'.

We call $\mathrel{\Box\!\!\rightarrow}$ the counterfactual (or subjunctive) conditional operator.

Then:[6]

Causal expected utility

$$\mathrm{CEU}(O) = \sum_{w} p(O \mathrel{\Box\!\!\rightarrow} w) \cdot u(w)$$

In a slogan, CDT thus says: choose the option that you expect would have the best outcome, were you to choose it.

[6] Exercise: give the "states" formulation of causal expected utility (CEU), instead of the "worlds" formulation.

You will notice that we are appealing to counterfactuals here, when we motivated two-boxing via *causal* considerations. What gives?

Counterfactuals and causation

You will notice that we are appealing to counterfactuals here, when we motivated two-boxing via *causal* considerations. What gives?

The idea is that, here, we're understanding the probabilities $p(O \mathrel{\Box\kern-1.5pt\rightarrow} S)$ as tracking the degree to which you believe your action would causally bring about the given state.

Counterfactuals and causation

You will notice that we are appealing to counterfactuals here, when we motivated two-boxing via *causal* considerations. What gives?

The idea is that, here, we're understanding the probabilities $p(O \mathrel{\Box\!\!\rightarrow} S)$ as tracking the degree to which you believe your action would causally bring about the given state.

One way to motivate this is to note the tight connection between counterfactuals and causation: When you want to know whether A caused C, for instance, it's natural to ask yourself questions like 'If A hadn't happened, would C have happened?'

Counterfactuals and causation

You will notice that we are appealing to counterfactuals here, when we motivated two-boxing via *causal* considerations. What gives?

The idea is that, here, we're understanding the probabilities $p(O \mathrel{\Box\!\!\rightarrow} S)$ as tracking the degree to which you believe your action would causally bring about the given state.

One way to motivate this is to note the tight connection between counterfactuals and causation: When you want to know whether A caused C, for instance, it's natural to ask yourself questions like 'If A hadn't happened, would C have happened?'

That said, we'll be spending a lot more time on the relationship between counterfactuals and causal notions tomorrow.

CDT in Newcomb

With counterfactuals understood in this causal way, it becomes clear that CDT recommends two-boxing in Newcomb.

CDT in Newcomb

With counterfactuals understood in this causal way, it becomes clear that CDT recommends two-boxing in Newcomb.

After all, since the contents of the opaque box are causally independent of what you do, your credence that choosing one box *would* result in getting a million is just equal to your credence that the million is there already.

So: $p(\mathit{One\ box} \mathrel{\Box\!\!\rightarrow} \mathit{Million}) = p(\mathit{Million})$. Likewise $p(\mathit{Both\ boxes} \mathrel{\Box\!\!\rightarrow} \mathit{Million}) = p(\mathit{Million})$. And so on, and so forth.

CDT in Newcomb

With counterfactuals understood in this causal way, it becomes clear that CDT recommends two-boxing in Newcomb.

After all, since the contents of the opaque box are causally independent of what you do, your credence that choosing one box *would* result in getting a million is just equal to your credence that the million is there already.

So: $p(\textit{One box} \,\Box\!\!\rightarrow \textit{Million}) = p(\textit{Million})$. Likewise $p(\textit{Both boxes} \,\Box\!\!\rightarrow \textit{Million}) = p(\textit{Million})$. And so on, and so forth.

	Million	*No million*
One-box	\$1,000,000	\$0
Two-box	\$1,001,000	\$1,000

CDT in Newcomb

With counterfactuals understood in this causal way, it becomes clear that CDT recommends two-boxing in Newcomb.

After all, since the contents of the opaque box are causally independent of what you do, your credence that choosing one box *would* result in getting a million is just equal to your credence that the million is there already.

So: $p(\textit{One box} \,\Box\!\!\rightarrow \textit{Million}) = p(\textit{Million})$. Likewise $p(\textit{Both boxes} \,\Box\!\!\rightarrow \textit{Million}) = p(\textit{Million})$. And so on, and so forth.

	Million	*No million*
One-box	\$1,000,000	\$0
Two-box	\$1,001,000	\$1,000

So, CDT calculates CEU in this case, just like ordinary EU theory. (*Pssst!* Here again, do we even need to do a calculation?)

Suppositional decision theories

CDT and EDT are different theories. But today, we want to close by showing that they have a common core—and by relating this fact to our study of *conditionals*.

The general recipe

CDT and EDT are different theories. But today, we want to close by showing that they have a common core—and by relating this fact to our study of *conditionals*.

Recall that EDT says: choose the option that you expect will have the best outcome, conditional on it being chosen.

The general recipe

CDT and EDT are different theories. But today, we want to close by showing that they have a common core—and by relating this fact to our study of *conditionals*.

Recall that EDT says: choose the option that you expect will have the best outcome, conditional on it being chosen.

In contrast, CDT says: choose the option that you expect would provide the best outcome, if it were to be chosen.

The general recipe

CDT and EDT are different theories. But today, we want to close by showing that they have a common core—and by relating this fact to our study of *conditionals*.

Recall that EDT says: choose the option that you expect will have the best outcome, conditional on it being chosen.

In contrast, CDT says: choose the option that you expect would provide the best outcome, if it were to be chosen.

In a sense, then, both theories say that you should evaluate options, by weighing how good their outcomes would be under the supposition that the options are chosen.

Supposition functions

We can make this idea precise by introducing the notation of a supposition function:[7]

$$s : \Delta(\mathcal{P}(\mathcal{W})) \times \mathcal{P}(\mathcal{W}) \to \Delta(\mathcal{P}(\mathcal{W})),$$

where $\Delta(\mathcal{P}(\mathcal{W}))$ is the set of all probability functions defined over $\mathcal{P}(\mathcal{W})$.

[7] We allow this function to be a mere partial function.

Supposition functions

We can make this idea precise by introducing the notation of a supposition function:[7]

$$s : \Delta(\mathcal{P}(\mathcal{W})) \times \mathcal{P}(\mathcal{W}) \to \Delta(\mathcal{P}(\mathcal{W})),$$

where $\Delta(\mathcal{P}(\mathcal{W}))$ is the set of all probability functions defined over $\mathcal{P}(\mathcal{W})$.

Thus, given a probability function p and proposition X, s returns a new probability function—intuitively, your probability function under the supposition of X, $s(p, X)$.

[7] We allow this function to be a mere partial function.

Supposition functions

We can make this idea precise by introducing the notation of a supposition function:[7]

$$s : \Delta(\mathcal{P}(\mathcal{W})) \times \mathcal{P}(\mathcal{W}) \to \Delta(\mathcal{P}(\mathcal{W})),$$

where $\Delta(\mathcal{P}(\mathcal{W}))$ is the set of all probability functions defined over $\mathcal{P}(\mathcal{W})$.

Thus, given a probability function p and proposition X, s returns a new probability function—intuitively, your probability function under the supposition of X, $s(p, X)$.

We make only two assumptions about this function:

- For all propositions A, $s(p, A)(A) = 1$.

[7] We allow this function to be a mere partial function.

We can make this idea precise by introducing the notation of a supposition function:[7]

$$s : \Delta(\mathcal{P}(\mathcal{W})) \times \mathcal{P}(\mathcal{W}) \to \Delta(\mathcal{P}(\mathcal{W})),$$

where $\Delta(\mathcal{P}(\mathcal{W}))$ is the set of all probability functions defined over $\mathcal{P}(\mathcal{W})$.

Thus, given a probability function p and proposition X, s returns a new probability function—intuitively, your probability function under the supposition of X, $s(p, X)$.

We make only two assumptions about this function:

- For all propositions A, $s(p, A)(A) = 1$.
- $s(p, \mathcal{W}) = p$

[7] We allow this function to be a mere partial function.

Indicative supposition

Suppositions come in different “modes”. For example, one kind of supposition is indicative supposition.

Indicative supposition

Suppositions come in different “modes”. For example, one kind of supposition is indicative supposition.

To paraphrase Bennett: ‘Suppose Shakespeare didn’t write *Hamlet*. Then, did someone else write *Hamlet*?

Indicative supposition

Suppositions come in different "modes". For example, one kind of supposition is indicative supposition.

To paraphrase Bennett: 'Suppose Shakespeare didn't write *Hamlet*. Then, did someone else write *Hamlet*?

More generally: Suppose A is, as a matter of fact, true. Then, how likely is C to be true?

Indicative supposition

Suppositions come in different “modes”. For example, one kind of supposition is indicative supposition.

To paraphrase Bennett: ‘Suppose Shakespeare didn’t write *Hamlet*. Then, did someone else write *Hamlet*?

More generally: Suppose A is, as a matter of fact, true. Then, how likely is C to be true?

It’s widely agreed that indicative supposition corresponds to conditionalization (Joyce, 1999; Bradley, 2017):

Conditionalization

$$s(p, A)(C) = p(C \mid A) = \frac{p(A \cap C)}{p(A)} \quad \text{if } p(A) > 0$$

Subjunctive supposition

Another mode of supposition is subjunctive supposition.

Subjunctive supposition

Another mode of supposition is subjunctive supposition.

Paraphrasing Bennett: 'Suppose Shakespeare hadn't written *Hamlet*. Then, would someone else have written *Hamlet*?

Subjunctive supposition

Another mode of supposition is subjunctive supposition.

Paraphrasing Bennett: 'Suppose Shakespeare hadn't written *Hamlet*. Then, would someone else have written *Hamlet*?

More generally: Suppose A were true. Then, would C also be true?

Another mode of supposition is subjunctive supposition.

Paraphrasing Bennett: 'Suppose Shakespeare hadn't written *Hamlet*. Then, would someone else have written *Hamlet*?

More generally: Suppose A were true. Then, would C also be true?

It's widely agreed that subjunctive supposition corresponds to something that Lewis (1976) calls imaging. Given some semantic assumptions about counterfactuals, however—assumptions which we'll discuss tomorrow—this is equivalent to the following:

Imaging

$$s(p, A)(C) = p^A(C) = p(A \mathrel{\Box\!\!\rightarrow} C)$$

The general decision rule

Given the notion of a supposition function, we can define an abstract, schematic form of expected utility:

Suppositional expected utility

$$\mathrm{SEU}_s(O) = \sum_w s(p, O)(w) \cdot u(w)$$

The general decision rule

Given the notion of a supposition function, we can define an abstract, schematic form of expected utility:

Suppositional expected utility

$$\mathrm{SEU}_s(O) = \sum_w s(p, O)(w) \cdot u(w)$$

Every suppositional decision theory says: choose an option that maximizes SEU (according to its specific supposition function).

The general decision rule

Given the notion of a supposition function, we can define an abstract, schematic form of expected utility:

Suppositional expected utility

$$\mathrm{SEU}_s(O) = \sum_w s(p, O)(w) \cdot u(w)$$

Every suppositional decision theory says: choose an option that maximizes SEU (according to its specific supposition function).

EDT and CDT fall out as special cases of this schema:

- EDT is the case where $s(p, O) = p(\cdot \mid O)$.

The general decision rule

Given the notion of a supposition function, we can define an abstract, schematic form of expected utility:

Suppositional expected utility

$$\mathrm{SEU}_s(O) = \sum_w s(p, O)(w) \cdot u(w)$$

Every suppositional decision theory says: choose an option that maximizes SEU (according to its specific supposition function).

EDT and CDT fall out as special cases of this schema:

- EDT is the case where $s(p, O) = p(\cdot \mid O)$.
- CDT is the case where $s(p, O) = p(O \mathbin{\Box\!\!\rightarrow} \cdot)$

Suppositions and conditionals

Why is this relevant?

Suppositions and conditionals

Why is this relevant?

Well, think about the following ways of evaluating an option:

- Suppose I choose O. Then, what will/would happen?

Why is this relevant?

Well, think about the following ways of evaluating an option:

- Suppose I choose O. Then, what will/would happen?
- If I choose O, what will/would happen?

Why is this relevant?

Well, think about the following ways of evaluating an option:

- Suppose I choose O. Then, what will/would happen?
- If I choose O, what will/would happen?

These seem to be saying the same thing. Thus, the debate about which suppositional decision theory is right can arguably be framed as a debate about the correct conditionals of deliberation...

Every supposition has a conditional representation

Representation result
For every supposition function s, we can define a family of (possibly credence-dependent) conditional operators $\{\Rightarrow_{s,p}\}_p$ by:

$$p(A \Rightarrow_{s,p} C) := s(p, A)(C),$$

whenever $s(p, A)$ is defined.

Every supposition has a conditional representation

Representation result
For every supposition function s, we can define a family of (possibly credence-dependent) conditional operators $\{\Rightarrow_{s,p}\}_p$ by:

$$p(A \Rightarrow_{s,p} C) := s(p, A)(C),$$

whenever $s(p, A)$ is defined.

Thus, every suppositional decision theory can be written in conditional form:

$$\mathrm{SEU}(O) = \sum_w p(O \Rightarrow_{s,p} w) \cdot u(w).$$

Every supposition has a conditional representation

Representation result
For every supposition function s, we can define a family of (possibly credence-dependent) conditional operators $\{\Rightarrow_{s,p}\}_p$ by:

$$p(A \Rightarrow_{s,p} C) := s(p, A)(C),$$

whenever $s(p, A)$ is defined.

Thus, every suppositional decision theory can be written in conditional form:

$$\mathrm{SEU}(O) = \sum_w p(O \Rightarrow_{s,p} w) \cdot u(w).$$

Qualification: this is a representation at the level of *credence*. It need not supply a single conditional proposition whose content is fixed independently of p. (This will be relevant when we discuss Lewisian triviality.)

Every supposition has a conditional representation

Representation result
For every supposition function s, we can define a family of (possibly credence-dependent) conditional operators $\{\Rightarrow_{s,p}\}_p$ by:

$$p(A \Rightarrow_{s,p} C) := s(p, A)(C),$$

whenever $s(p, A)$ is defined.

Thus, every suppositional decision theory can be written in conditional form:

$$\mathrm{SEU}(O) = \sum_{w} p(O \Rightarrow_{s,p} w) \cdot u(w).$$

Qualification: this is a representation at the level of *credence*. It need not supply a single conditional proposition whose content is fixed independently of p. (This will be relevant when we discuss Lewisian triviality.)

The choice of s is therefore the choice of a conditional of deliberation.

Summing up

Takeaways

What have we learned?

1. EU theory is natural, but assumes options and states are independent.

Takeaways

What have we learned?

1. EU theory is natural, but assumes options and states are independent.
2. EDT handles act-state dependence by conditionalizing on the option.

Takeaways

What have we learned?

1. EU theory is natural, but assumes options and states are independent.
2. EDT handles act-state dependence by conditionalizing on the option.
3. Newcomb pressures EDT because evidential and causal dependence come apart.

Takeaways

What have we learned?

1. EU theory is natural, but assumes options and states are independent.
2. EDT handles act-state dependence by conditionalizing on the option.
3. Newcomb pressures EDT because evidential and causal dependence come apart.
4. Stalnaker-style CDT builds counterfactuals directly into expected utility.

Takeaways

What have we learned?

1. EU theory is natural, but assumes options and states are independent.
2. EDT handles act-state dependence by conditionalizing on the option.
3. Newcomb pressures EDT because evidential and causal dependence come apart.
4. Stalnaker-style CDT builds counterfactuals directly into expected utility.
5. EDT and CDT are both instances of a general suppositional decision rule—one which can itself be cashed out in the language of conditionals.

Takeaways

What have we learned?

1. EU theory is natural, but assumes options and states are independent.
2. EDT handles act-state dependence by conditionalizing on the option.
3. Newcomb pressures EDT because evidential and causal dependence come apart.
4. Stalnaker-style CDT builds counterfactuals directly into expected utility.
5. EDT and CDT are both instances of a general suppositional decision rule—one which can itself be cashed out in the language of conditionals.

Tomorrow, we'll start thinking more about the precise semantics (etc.) for the conditionals of deliberation, which feature in theories like CDT.

References

Jonathan Bennett. Farewell to the phlogiston theory of conditionals. *Mind*, XCVII(388):509–527, 1988. ISSN 1460-2113. doi: 10.1093/mind/xcvii.388.509. URL http://dx.doi.org/10.1093/mind/xcvii.388.509.

Richard Bradley. *Decision Theory with a Human Face*. Cambridge University Press, 2017.

K. DeRose. The conditionals of deliberation. *Mind*, 119(473):1–42, jan 2010. doi: 10.1093/mind/fzp149. URL https://doi.org/10.1093%2Fmind%2Ffzp149.

Allan Gibbard and William Harper. Counterfactuals and two kinds of expected utility. In A. Hooker, J. J. Leach, and E. F. McClennen, editors, *Foundations and Applications of Decision Theory*, pages 125–162. D. Reidel, 1978.

Richard C. Jeffrey. *The Logic of Decision*. 2nd Edition, Chicago and London: University of Chicago Press, 1983.

James M. Joyce. *The Foundations of Causal Decision Theory*. Cambridge University Press, 1999.

David Lewis. Probabilities of conditionals and conditional probabilities. *The Philosophical Review*, 85(3):297, Jul 1976. ISSN 0031-8108. doi: 10.2307/2184045. URL http://dx.doi.org/10.2307/2184045.

David Lewis. 'why ain'cha rich?'. *Noûs*, 15(3):377, Sep 1981. ISSN 0029-4624. doi: 10.2307/2215439. URL http://dx.doi.org/10.2307/2215439.

References ii

Robert Nozick. Newcomb's problem and two principles of choice. In Nicholas Rescher, editor, *Essays in Honor of Carl G. Hempel*, pages 114–146. Reidel, 1969.

Robert Stalnaker. Letter to David Lewis. In William Harper, Robert C. Stalnaker, and Glenn Pearce, editors, *Ifs*, pages 151–152. Reidel, 1981.

Timothy Williamson. *Suppose and Tell: The Semantics and Heuristics of Conditionals*. Oxford, England: Oxford University Press, 2020.