Decision-making under Determinism

Conditionals in Decision Theory
Lecture 3

Calum McNamara and Paolo Santorio
IU and UMD

ESSLLI 2026, Prague

Similarity for subjunctives

(1) If I had struck this match, it would have lit.

- One might hope that an intuitive notion of similarity will be sufficient.
- But this is shown not to be true by the Nixon case

(2) If Nixon had pushed the button [that launches nuclear missiles], there would have been a nuclear holocaust.

Intuitively true, and (if similarity is naïve similarity) predicted to be false.

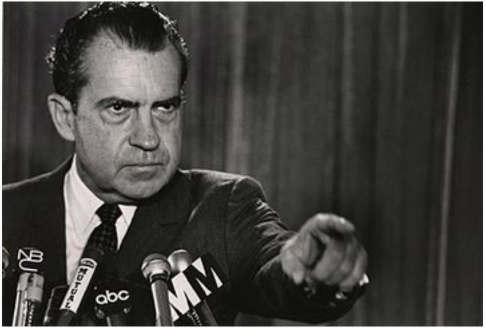

The counterfactual "If Nixon had pressed the button there would have been a nuclear holocaust" is true or can be imagined to be so. Now suppose that there never will be a nuclear holocaust. Then that counterfactual is, [on an analysis that employs a naïve similarity relation], very likely false. For given any world in which antecedent and consequent are both true it will be easy to imagine a closer world in which the antecedent is true and the consequent false. For we need only imagine a change that prevents the holocaust but that does not require such a great divergence from reality. (Fine 1975)

The nature of similarity

The solution: similarity exploits a past-future asymmetry.

Two main criteria for similarity:

- match in history up to the reference time of the antecedent;
- match with actual laws of nature.

Unfortunately: given minimal assumptions, we cannot assume that similarity requires both perfect match in history and perfect match with actual laws.

The nature of similarity

- Let L_w be the proposition that encodes all the laws of nature in w;
- Let $H_{w,t}$ be the proposition that encodes all particular historical events in history in w, up to time t.
- Let H_w be the proposition that encodes all of the history of w.

Let us consider the hypothesis of determinism, defined as follows:

Determinism. L_w and $H_{w,t}$ jointly entail H_w (for any t).

- If Determinism is true, then worlds that we reach via counterfactual supposition cannot verify both actual history and actual laws.
  (At least, unless we want to make all counterfactuals trivially true.)
- Perhaps determinism is false, but we want our semantics for counterfactuals to be general enough to apply to both the deterministic and the indeterministic case.

Given this: we have two options.

The nature of similarity

Option #1: miracles. This is the option adopted by Lewis.

Lewis's criteria for similarity:

(1) It is of the first importance to avoid big, widespread, diverse violations of law.

(2) It is of the second importance to maximize the spatiotemporal region throughout which perfect match of particular fact prevails.

(3) It is of the third importance to avoid even small, localized, simple violations of law.

(4) It is of little or no importance to secure approximate similarity of particular fact, even in matters that concern us greatly.

So Lewis chooses:

(i) perfect match in (past) history;

(ii) (possibly) imperfect match in laws.

Most similar -worlds to w: worlds with an identical course of history up to shortly before t, and where a 'small miracle' brings about .

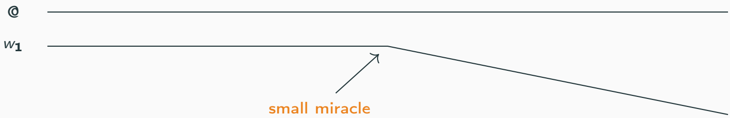

In the Nixon case: closest worlds are worlds where:

- History matches actual history until shortly before antecedent time;
- A 'small miracle' leads Nixon to push the button.
- History takes a very different course after that.

Hence: it's predicted to be true that, if Nixon had pressed the button, there would have been a nuclear holocaust.

The nature of similarity

In the Nixon case: closest worlds are worlds where:

- History matches actual history until shortly before antecedent time;
- A 'small miracle' leads Nixon to push the button.
- History takes a very different course after that.

Hence: it's predicted to be true that, if Nixon had pressed the button, there would have been a nuclear holocaust.

A peculiarity of the account: conditionals of this sort are predicted to be true:

(4) If I had scratched my head once more yesterday, the laws would have been different.

The nature of similarity

Option #2: no-miracles. This option is defended by a number of philosophers, including Dorr (2016) a.o.

On this option, we require:

- Laws of nature hold throughout closest worlds.
- History is so-very-slightly different, starting from the beginning of time.

So we choose:

(i) perfect match in laws;

(ii) (possibly) imperfect match in history.

Most similar -worlds to w: worlds with identical laws to w, and with a history that is different from the very beginning (though, possibly, in imperceptible ways).

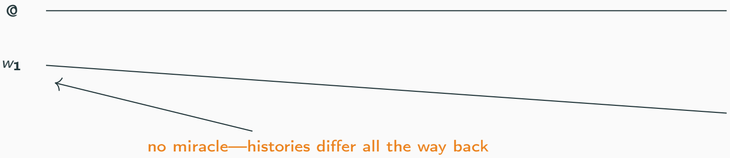

The nature of similarity

In the Nixon case: closest worlds are worlds where:

- History was always different, from the beginning of time.
- This divergence does not generate any divergence in macroscopic facts, up to the point at which Nixon to push the button.
- History takes a very different course after that.

Again: it's predicted to be true that, if Nixon had pressed the button, there would have been a nuclear holocaust.

The nature of similarity

In the Nixon case: closest worlds are worlds where:

- History was always different, from the beginning of time.
- This divergence does not generate any divergence in macroscopic facts, up to the point at which Nixon to push the button.
- History takes a very different course after that.

Again: it's predicted to be true that, if Nixon had pressed the button, there would have been a nuclear holocaust.

A peculiarity of the account: conditionals of this sort are predicted to be true:

(6) If I had scratched my head once more yesterday, the whole history of the world starting from the Big Bang would have been different.

Today we will discuss:

1. Why Lewis's "miracles account" of similarity for counterfactuals initially looks fit for purpose for CDT.
2. Arif Ahmed's *Betting on the Laws* case—an apparent counterexample to CDT.

What conditionals for deliberation?

DeRose on conditionals of deliberation

So now that we have truth conditions, we can ask: what conditionals are employed in deliberation—indicative, subjunctive, or both?

- Some arguments obviously come from decision-theoretic considerations (such as: Newcomb-type cases).
- But some other arguments can be made directly from the language of deliberation.

One argument we consider today: DeRose (2010).

DeRose on conditionals of deliberation

DeRose (2010) considers (variants of) the following:

(7) a. If I drive to work without filling up the tank, I will run out of petrol.
    b. If I leave for my meeting 5 minutes before it starts, I will be late.
    c. If the house is not painted, it will soon look quite shabby.

> *These all seem to be conditionals that would be useful in deliberation: to the extent that I have reason to believe one of them, then in so far as I desire its consequent to be true, I have reason to make (or to try to make, in cases where the truth of the antecedent is not completely up to me) its antecedent true.* *(DeRose 2010)*

DeRose on conditionals of deliberation

(7) a. If I drive to work without filling up the tank, I will run out of petrol.
  b. If I leave for my meeting 5 minutes before it starts, I will be late.
  c. If the house is not painted, it will soon look quite shabby.

(P1) These conditionals are paradigm examples of conditionals of deliberation.

(P2) These conditionals are not subjunctive, but rather indicative.

(C) So conditionals of deliberation are indicative and not subjunctive—contrary to what causal decision theorists suggest.

How good is this argument?

The logical form of *will*-conditionals

To answer this: we need to say more about two issues:

- the meaning of *will*;
- the compositional semantics of conditionals.

The meaning of *will*.

- Traditionally, especially in tense logic, *will* is treated as a tense.
- But there are excellent arguments that *will* is actually a modal auxiliary.

We will consider two arguments for the claim that *will* is a modal.

Argument #1: modal subordination.

It is well known that modal expressions can enter anaphoric relations.

Here is a classical example, due to Craige Roberts:

(8) A wolf might come in. It would eat you. (Roberts 1989)

- The *would*-claim is understood as a conditional: it says that the wolf would eat the interlocutor *in the scenario in which the wolf comes in*.
- Hence the second sentence in the discourse is subordinated, in the relevant sense, to the first sentence.

Crucially, as Klecha (2014) points out, *will* displays an analogous behavior:

(9) a. If Katie travels to Berkeley, she will shop at Amoeba Records. She will buy a boxed set and a dozen used LPs. (Cariani 2021)

b. The supplies might arrive tomorrow. It will be late in the day. (Cariani and Santorio 2018)

c. Don't go near that bomb! It'll explode. (Binnick 1971)

Parallel sentences involving past tense cannot be forced to undergo subordination. Consider:

(10) a. If Katie traveled to Berkeley, she shopped at Amoeba Records. She bought a boxed set and a dozen used LPs.

b. The supplies might have arrived yesterday. It was late in the day.

c. I hope he didn't go near that bomb! It exploded.

Will: motivating a modal analysis

Argument #2: the acquaintance inference.

Background: Predicates of personal taste, like *tasty*, *delicious*, and *good* and *great* in some uses, trigger an inference to the effect that the speaker has had a certain relevant kind of experience.

To see this, notice that the continuations in (11) are all pragmatically odd.

(11) a. That sandwich is delicious, # but I haven't tried it.

b. Joan's book is great, # but I haven't read it.

c. The soup is disgusting, # but I haven't had it.

Will: motivating a modal analysis

Interestingly, this effect is suppressed under modals, including strong epistemic modals like *must* and *have to*.

(12) a. That sandwich must/has to/should/... be delicious, but I haven't tried it.

b. Joan's book must/has to/should/... to be great, but I haven't read it.

c. The soup must/has to/should/... be disgusting, but I haven't had it.

Will patterns with modals.

(13) a. That sandwich will be delicious, but I haven't tried it/won't try it.

b. Joan's book will be great, but I haven't read it/won't read it.

c. The soup will be disgusting, but I haven't had it/will not have it.

A modal analysis for *will*

If *will* is a modal, what analysis should we adopt?

The obvious suggestion: an analysis on which it quantifies or selects from metaphysical possibilities, similarly to subjunctive *would*.
(Copley 2009, Cariani and Santorio 2018).

My favorite option: a selection semantics, which makes it very close to the semantics for *would*.

(14) $[w, f, s]$*will* A= true iff $[s(w, s(f(w)), f, s]A$= true

The compositional semantics of conditionals.

- So far, we have presupposed that, at the level of logical form, conditionals simply involve a binary connective.

$$\textit{match struck} \;\Box\!\!\rightarrow\; \textit{match lights}$$

- But this is not the standard view of conditionals in linguistic semantics.
- On the standard view, the restrictor theory (Kratzer 1981, 1986, 2012), conditionals (i) involve a modal auxiliary and (ii) *if*-clauses are used to restrict the modal.

(15) a. If Alice is in Prague, Bob must be in Prague too.

b. must [if Alice is in Prague] [Bob is in Prague]

(16) a. If Alice is in Prague, Bob must be in Prague.

b.

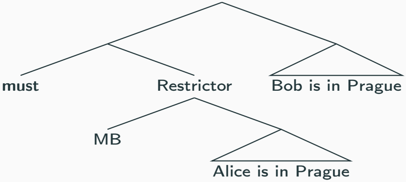

- This doesn't make a difference at the level of truth conditions.
- Where it matters: predicting how conditionals and overt modals interact.

What about conditionals that do not include an overt modal (bare conditionals)?

(17) If Maria danced, Frida danced.

- The standard view: they include a covert epistemic modal.
- So e.g. the logical form of (17) is:

(18) MUST [if Maria danced] [Frida danced]

The LF of *will*-conditionals

Given this treatment of *will* and this LF for conditionals, we can make some observations about *will*-conditionals (following Ciardelli and Ommundsen 2025).

(19) If you strike this match, it will light.

We predict that (19) has two possible logical forms:

- single modal:

  will [you strike the match][match lights]
- double modal:

  MUST [you strike the match][*will* [match lights]]

Interestingly, these two conditionals will differ in truth conditions!

An ambiguity in *will*-conditionals

(19) If you strike this match, it will light.

Single modal:

will [you strike the match][match lights]

- *will* is the main modal of the conditional;
- the conditional has causal/metaphysical flavor.

Double modal:

MUST [you strike the match][*will* [match lights]]

- covert MUST is the main modal of the conditional;
- the conditional has epistemic flavor.

So the prediction is: systematically, *will*-conditionals have two readings!

So: on one reading, *will*-conditionals are almost synonymous with subjunctives!

Indeed, this is what we have observed:

(20) a. If I strike this match, it will light.

b. If I were to strike this match, it would light.

- The only difference is that (20b) carries a 'remoteness' inference of sorts (Iatridou 2000).
- What this is is a matter for another day.
- The key thing: both (20a) and (20b) have causal/metaphysical flavor.

An ambiguity in *will*-conditionals

Back to DeRose's observation:

(7) a. If I drive to work without filling up the tank, I will run out of petrol.
  b. If I leave for my meeting 5 minutes before it starts, I will be late.
  c. If the house is not painted, it will soon look quite shabby.

- DeRose is correct that conditionals in (7) are used in deliberation.
- But we can explain why without concluding that indicatives in general are conditionals of deliberation.
- We know that the conditionals in (7) are systematically ambiguous.
- They are used in deliberation when they receive their causal/metaphysical reading (i.e. the same flavor we find in subjunctives).

An ambiguity in *will*-conditionals

A general takeaway: the classical philosophical distinction between indicatives and subjunctives is way too crude.

A better taxonomy (to be still refined) exploits two dimensions:

- modal flavor (epistemic vs metaphysical);
- remoteness (exactly what this is... we'll see in another ESSLLI class!).

	No remoteness effect	Remoteness effect
Epistemic	Bare cond, epistemic *will*-cond *If Frida danced, Maria danced* MUST *[if I strike] [will [it lights]]*	[empty?]
Metaphysical	metaphysical will-conditionals *will [if I strike] [it lights]*	*would*-conditionals *if I struck, it would light*

CDT and the miracles account

Stalnakerian CDT, revisited

Recall the version of CDT, Stalnakerian CDT, we entertained on Day 1:

Causal expected utility

$$\mathrm{CEU}(O) = \sum_i p(O \mathrel{\Box\!\!\rightarrow} S_i) \cdot u(O \cap S_i)$$

Informally: Choose the option that you expect *would* have the best result, *were* you to choose it.

If this theory is going to work as intended, then we need the counterfactuals in it to hold fixed things that are not causally influenced by your choices.

Why this vindicates two-boxing

The miracles account seems to vindicate this.

To see this, recall In Newcomb's problem:

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

The miracles account holds fixed the box's contents, when we suppose that you make a different choice today (because it holds history fixed).

Upshot: two-boxing does better than one-boxing, by dominance reasoning.

Causal independence

Lewis's miracles account looks good for CDT, because it holds the past fixed; and we typically think of the past as being causally independent of what you do.

However, Lewis's account does not hold fixed the laws of nature. The closest world at which you do otherwise than what you actually do is one at which the laws are different.

But laws of nature, we usually think, are *also* causally independent of what you do. What gives?

The miracles account doesn't hold fixed *everything* that's causally independent of your actions.

Causal independence as similarity?

The miracles account doesn't hold fixed *everything* that's causally independent of your actions.

Arguably, then, it's not perfectly well equipped for use in CDT.

Causal independence as similarity?

The miracles account doesn't hold fixed *everything* that's causally independent of your actions.

Arguably, then, it's not perfectly well equipped for use in CDT.

Can't we, then, say that the "most similar world" is one that holds fixed *everything* that's causally independent of your actions?

Causal independence as similarity

No. Suppose that a moment ago, you raised your arm.

Now, suppose that the most similar world at which you *didn't* raise your arm is one that holds fixed *everything* that's causally independent of your arm-raising.

This includes (i) the complete history of the world (since there's no backwards causation), and also (ii) the laws of nature.

But if the laws are deterministic, the history of the world and the laws jointly determine *whatever you do*. So, they must determine that you *did* raise your arm. So a contradiction is true.

Upshot: under determinism, we can't hold fixed everything that's causally independent of your choice.[1]

[1] Or can we? We'll discuss this possibility at the end of the lecture.

CDT under determinism

Maybe the miracles account is as close as we can get to holding everything that's causally independent of your choice fixed.

Matters arising

Maybe the miracles account is as close as we can get to holding everything that's causally independent of your choice fixed.

At the very least, it gets Newcomb's problem right.

Maybe the miracles account is as close as we can get to holding everything that's causally independent of your choice fixed.

At the very least, it gets Newcomb's problem right.

However, Arif Ahmed (2013) has given a case which shows that, if CDT'ers adopt the miracles account, then they seem to face a devastating counterexample.

A lemma

But first, a lesson in probability (and logic).

A lemma

But first, a lesson in probability (and logic).

Recall the truth table for the *material conditional* $A \supset C$:

A	C	$A \supset C$
T	T	T
T	F	F
F	T	T
F	F	T

A	C	$\neg A \vee C$
T	T	T
T	F	F
F	T	T
F	F	T

$A \wedge C$	$A \wedge \neg C$
$\neg A \wedge C$	$\neg A \wedge \neg C$

A lemma

Lemma: Suppose $p(A \supset C) = 1$. Then

$$p(A) \leq p(C).$$

A lemma

Lemma: Suppose $p(A \supset C) = 1$. Then

$$p(A) \leq p(C).$$

Proof: Since $A \supset C$ is false only when $A \wedge \neg C$ is true:

$$p(A \wedge \neg C) = 0.$$

A lemma

Lemma: Suppose $p(A \supset C) = 1$. Then

$$p(A) \leq p(C).$$

Proof: Since $A \supset C$ is false only when $A \wedge \neg C$ is true:

$$p(A \wedge \neg C) = 0.$$

Therefore:

$$p(A) = p(A \wedge C) + p(A \wedge \neg C) = p(A \wedge C) \leq p(C).$$

Now, here is Ahmed's case (Ahmed, 2013):

> Betting on the Laws. *Imagine there's a proposition L, describing the (deterministic) laws of nature, in which you are almost certain. So, $p(L) \approx 1$. (Maybe, for example, L describes some particular deterministic formulation of quantum mechanics—like Bohmian mechanics or the Everett interpretation.) Now, imagine I offer you the following bets to choose between—you must choose one of them. First, there's Bet 1, which pays \$1 if L is true, but pays nothing if L is false. Second, there's Bet 2, which pays nothing if L is true, but pays \$1 if L is false. You're certain that nothing you do can causally affect L's truth-value, and you care only about winning the dollar.*

Here is the decision table:

	L	$\neg L$
Bet 1	\$1	\$0
Bet 2	\$0	\$1

Betting on the Laws (cont'd)

Here is the decision table:

	L	$\neg L$
Bet 1	\$1	\$0
Bet 2	\$0	\$1

Question: Which option should you choose?

What must CDT compare?

Intuitively, you should choose Bet 1.

What must CDT compare?

Intuitively, you should choose Bet 1.

After all, you are almost certain that L is true, and certain that your choice cannot causally affect whether L is true.

What must CDT compare?

Intuitively, you should choose Bet 1.

After all, you are almost certain that L is true, and certain that your choice cannot causally affect whether L is true.

But what does CDT say?

What must CDT compare?

Intuitively, you should choose Bet 1.

After all, you are almost certain that L is true, and certain that your choice cannot causally affect whether L is true.

But what does CDT say?

Notice that:

$$\begin{aligned}
\text{CEU}(\textit{Bet 1}) &= p(\textit{Bet 1} \mathrel{\Box\!\!\rightarrow} L) \cdot 1 + p(\textit{Bet 1} \mathrel{\Box\!\!\rightarrow} \neg L) \cdot 0 \\
&= p(\textit{Bet 1} \mathrel{\Box\!\!\rightarrow} L) \\
\text{CEU}(\textit{Bet 2}) &= p(\textit{Bet 2} \mathrel{\Box\!\!\rightarrow} L) \cdot 0 + p(\textit{Bet 2} \mathrel{\Box\!\!\rightarrow} \neg L) \cdot 1 \\
&= p(\textit{Bet 2} \mathrel{\Box\!\!\rightarrow} \neg L)
\end{aligned}$$

What must CDT compare?

Intuitively, you should choose Bet 1.

After all, you are almost certain that L is true, and certain that your choice cannot causally affect whether L is true.

But what does CDT say?

Notice that:

$$\begin{aligned} \mathrm{CEU}(\textit{Bet 1}) &= p(\textit{Bet 1} \mathrel{\Box\!\!\rightarrow} L) \cdot 1 + p(\textit{Bet 1} \mathrel{\Box\!\!\rightarrow} \neg L) \cdot 0 \\ &= p(\textit{Bet 1} \mathrel{\Box\!\!\rightarrow} L) \\ \mathrm{CEU}(\textit{Bet 2}) &= p(\textit{Bet 2} \mathrel{\Box\!\!\rightarrow} L) \cdot 0 + p(\textit{Bet 2} \mathrel{\Box\!\!\rightarrow} \neg L) \cdot 1 \\ &= p(\textit{Bet 2} \mathrel{\Box\!\!\rightarrow} \neg L) \end{aligned}$$

So, CDT's answer entirely depends on what credences you assign to the counterfactuals $\textit{Bet 1} \mathrel{\Box\!\!\rightarrow} L$ and $\textit{Bet 2} \mathrel{\Box\!\!\rightarrow} \neg L$, respectively.

A derivation

So what are those credences?

A derivation

So what are those credences?

Well, take an arbitrary world w, and suppose:

$$\textit{Bet 1} \mathrel{\Box\kern-1.5pt\rightarrow} L$$

is true at w.

A derivation

So what are those credences?

Well, take an arbitrary world w, and suppose:

$$\textit{Bet 1} \mathrel{\Box\!\!\rightarrow} L$$

is true at w.

At w, you must choose exactly one of the two bets, Bet 1 and Bet 2. Thus, at w:

$$\textit{Bet 1} \lor \textit{Bet 2}.$$

We now argue by cases.

First horn: you choose *Bet 1*

First suppose *Bet 1* is true at w.

First horn: you choose *Bet 1*

First suppose *Bet 1* is true at w.

Now, since *Bet 1* $\square\!\!\rightarrow L$ is also true at w, it follows by Modus Ponens that L is true at w.

First horn: you choose *Bet 1*

First suppose *Bet 1* is true at w.

Now, since *Bet 1* $\Box\!\!\rightarrow$ L is also true at w, it follows by Modus Ponens that L is true at w.

But what is the closest Bet 2-world?

First horn: you choose *Bet 1*

First suppose *Bet 1* is true at w.

Now, since *Bet 1* $\Box\!\!\rightarrow L$ is also true at w, it follows by Modus Ponens that L is true at w.

But what is the closest Bet 2-world?

By the miracles account, this must be a world that shares w's history, but *not* it's laws.

First suppose *Bet 1* is true at w.

Now, since *Bet 1* $\Box\!\!\rightarrow L$ is also true at w, it follows by Modus Ponens that L is true at w.

But what is the closest Bet 2-world?

By the miracles account, this must be a world that shares w's history, but *not* it's laws.

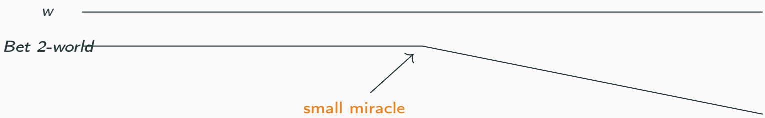

First horn: you choose *Bet 1*

First suppose *Bet 1* is true at w.

Now, since *Bet 1* $\Box\!\!\rightarrow L$ is also true at w, it follows by Modus Ponens that L is true at w.

But what is the closest Bet 2-world?

By the miracles account, this must be a world that shares w's history, but *not* it's laws.

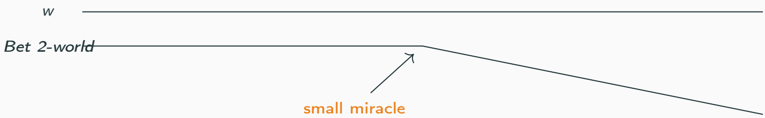

So, the closest *Bet*2-world to w is one at which $\neg L$ is true, and thus *Bet 2* $\Box\!\!\rightarrow \neg L$ is true at w.

Second horn: you choose *Bet 2*

Now suppose you choose Bet 2 at w instead.

Second horn: you choose *Bet 2*

Now suppose you choose Bet 2 at w instead.

By assumption, $\textit{Bet 1} \mathrel{\Box\!\!\rightarrow} L$ is true at w

Second horn: you choose *Bet 2*

Now suppose you choose Bet 2 at w instead.

By assumption, *Bet 1* $\Box\!\!\rightarrow L$ is true at w

But this means that the most similar *Bet 1*-world to w is an L-world (by our semantics).

Second horn: you choose *Bet 2*

Now suppose you choose Bet 2 at w instead.

By assumption, *Bet 1* $\boxright$ L is true at w

But this means that the most similar *Bet 1*-world to w is an L-world (by our semantics).

The most similar *Bet 1*-world to w shares the same history as w (by the miracles account).

Second horn: you choose *Bet 2*

Now suppose you choose Bet 2 at w instead.

By assumption, *Bet 1* $\Box\!\!\rightarrow$ L is true at w

But this means that the most similar *Bet 1*-world to w is an L-world (by our semantics).

The most similar *Bet 1*-world to w shares the same history as w (by the miracles account).

But because L is deterministic, this means w cannot also be a world at which L is true.

Second horn: you choose *Bet 2*

Now suppose you choose Bet 2 at w instead.

By assumption, $\textit{Bet 1} \mathrel{\Box\!\!\rightarrow} L$ is true at w

But this means that the most similar *Bet 1*-world to w is an L-world (by our semantics).

The most similar *Bet 1*-world to w shares the same history as w (by the miracles account).

But because L is deterministic, this means w cannot also be a world at which L is true.

So $\textit{Bet 2} \wedge \neg L$ is true at w, and thus $\textit{Bet 2} \mathrel{\Box\!\!\rightarrow} \neg L$ is true at w, too.

Closing the dilemma

We assumed that *Bet 1* $\Box\!\!\rightarrow$ *L* was true at an arbitrary world *w*.

Closing the dilemma

We assumed that *Bet 1* $\Box\!\!\rightarrow$ *L* was true at an arbitrary world w.

Also, we appealed to the fact that you must choose either Bet 1 or Bet 2: *Bet 1* $\vee$ *Bet 2*.

Closing the dilemma

We assumed that $\textit{Bet 1} \mathrel{\Box\!\!\rightarrow} L$ was true at an arbitrary world w.

Also, we appealed to the fact that you must choose either Bet 1 or Bet 2: $\textit{Bet 1} \vee \textit{Bet 2}$.

We then showed that, assuming *either* that you choose Bet 1 or Bet 2, we can derive the counterfactual $\textit{Bet 2} \mathrel{\Box\!\!\rightarrow} L$.

Since w was arbitrary, we therfore have proved the following material conditional:

The key implication

$$(\textit{Bet 1} \mathrel{\Box\!\!\rightarrow} L) \supset (\textit{Bet 2} \mathrel{\Box\!\!\rightarrow} \neg L)$$

From the dilemma to CDT's verdict

Having reasoned your way to this conclusion, you should therefore be certain of the material conditional:

$$p\big((\mathit{Bet\ 1} \mathrel{\Box\!\!\rightarrow} L) \supset (\mathit{Bet\ 2} \mathrel{\Box\!\!\rightarrow} \neg L)\big) = 1$$

Having reasoned your way to this conclusion, you should therefore be certain of the material conditional:

$$p\big((\mathit{Bet\ 1} \;\Box\!\!\rightarrow L) \supset (\mathit{Bet\ 2} \;\Box\!\!\rightarrow \neg L)\big) = 1$$

By now recall our lemma about the probabilities of antecedents and consequents of certain material conditionals:

$$p(\mathit{Bet\ 1} \;\Box\!\!\rightarrow L) \leq p(\mathit{Bet\ 2} \;\Box\!\!\rightarrow \neg L)$$

From the dilemma to CDT's verdict

Having reasoned your way to this conclusion, you should therefore be certain of the material conditional:

$$p\big((\textit{Bet 1} \mathrel{\Box\!\!\rightarrow} L) \supset (\textit{Bet 2} \mathrel{\Box\!\!\rightarrow} \neg L)\big) = 1$$

By now recall our lemma about the probabilities of antecedents and consequents of certain material conditionals:

$$p(\textit{Bet 1} \mathrel{\Box\!\!\rightarrow} L) \leq p(\textit{Bet 2} \mathrel{\Box\!\!\rightarrow} \neg L)$$

But these are the probabilities that the two bets would win:

$$\mathrm{CEU}(\textit{Bet 1}) = p(\textit{Bet 1} \mathrel{\Box\!\!\rightarrow} L), \qquad \mathrm{CEU}(\textit{Bet 2}) = p(\textit{Bet 2} \mathrel{\Box\!\!\rightarrow} \neg L).$$

From the dilemma to CDT's verdict

Having reasoned your way to this conclusion, you should therefore be certain of the material conditional:

$$p\big((\textit{Bet 1} \mathrel{\Box\!\!\rightarrow} L) \supset (\textit{Bet 2} \mathrel{\Box\!\!\rightarrow} \neg L)\big) = 1$$

By now recall our lemma about the probabilities of antecedents and consequents of certain material conditionals:

$$p(\textit{Bet 1} \mathrel{\Box\!\!\rightarrow} L) \leq p(\textit{Bet 2} \mathrel{\Box\!\!\rightarrow} \neg L)$$

But these are the probabilities that the two bets would win:

$$\mathrm{CEU}(\textit{Bet 1}) = p(\textit{Bet 1} \mathrel{\Box\!\!\rightarrow} L), \qquad \mathrm{CEU}(\textit{Bet 2}) = p(\textit{Bet 2} \mathrel{\Box\!\!\rightarrow} \neg L).$$

Therefore:

$$\boxed{\mathrm{CEU}(\textit{Bet 1}) \leq \mathrm{CEU}(\textit{Bet 2})}$$

So CDT says that choosing *Bet 2* is at least rationally permissible.

References

Arif Ahmed. Causal decision theory: A counterexample. *Philosophical Review*, 122(2): 289–306, 2013. doi: 10.1215/00318108-1963725. URL https://doi.org/10.1215%2F00318108-1963725.