Conditionals: between language and reasoning

Class 1 - Introduction

www.ivanociardelli.altervista.org/conditionals
ivano.ciardelli@lmu.de

Two switches

Imagine a long hallway with a light in the middle and with two switches, one at each end. One switch is called switch A and the other one is called switch B. As the following wiring diagram shows, the light is on whenever both switches are in the same position (both up or both down); otherwise, the light is off. Right now, switch A and switch B are both up, and the light is on. But things could be different. . .

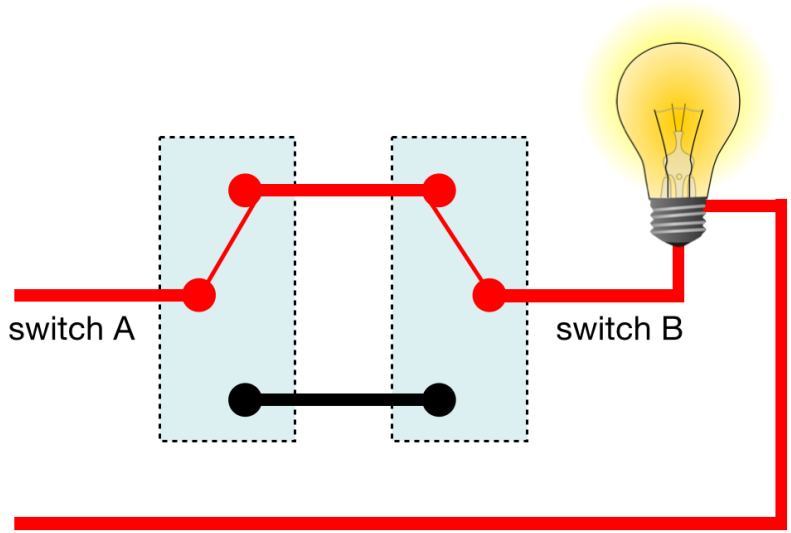

Firing squad

- A squad consisting of two riflemen, A and B, is about to shoot a prisoner.
- If the court orders the execution, the captain will give a signal.
- Each rifleman is posed to fire if, and only if, the captain signals.
- Either shot by itself is enough to kill the prisoner.

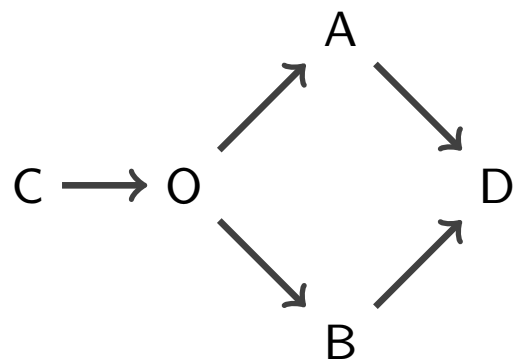

- As a matter of fact, the court does order the execution, so that the captain signals, the executioners shoot, and the prisoner dies.
- What if things had gone otherwise?

What is a conditional?

- A conditional is a sentence of the form if A, C, where A (the antecedent) and C (the consequent) are sentential clauses:

(1) a. If kangaroos had no tails, they would topple over.
b. If kangaroos have no tails, they've been fooling us all this time.

What is a conditional?

- A conditional is a sentence of the form if A, C, where A (the antecedent) and C (the consequent) are sentential clauses:

  (1) a. If kangaroos had no tails, they would topple over.
  b. If kangaroos have no tails, they've been fooling us all this time.

- This is a bit too narrow. There are other ways to express conditionals:

  (2) a. Had I known about the strike, I would have stayed home.
  b. No Martini, no party.
  c. Pay him and he'll tell you everything.
  d. In case of rain, the event will be canceled.

What is a conditional?

- A conditional is a sentence of the form if A, C, where A (the antecedent) and C (the consequent) are sentential clauses:

  (1) a. If kangaroos had no tails, they would topple over.
  b. If kangaroos have no tails, they've been fooling us all this time.

- This is a bit too narrow. There are other ways to express conditionals:

  (2) a. Had I known about the strike, I would have stayed home.
  b. No Martini, no party.
  c. Pay him and he'll tell you everything.
  d. In case of rain, the event will be canceled.

- The common core of these is the fact that they involve:
  - the creation of a hypothetical context/situation where A holds;
  - the claim that C holds in this hypothetical context.

- Most literature focuses on conditional statements.

- However, conditionalization is a general phenomenon, which applies to questions and imperatives as well:

  a. If we invite Alice, how will Bob react?
  b. If we had invited Alice, how would Bob have reacted?
  c. If we invite Alice, don't tell Bob that we did.

The importance of conditionals

Conditionals are probably the most widely studied of all linguistic/logical constructions. Why is that?

The importance of conditionals

Conditionals are probably the most widely studied of all linguistic/logical constructions. Why is that?

- Because they are difficult!

The importance of conditionals

Conditionals are probably the most widely studied of all linguistic/logical constructions. Why is that?

- Because they are difficult!
- Because conditionals play a key role in connection with a wide range of disciplines:
  - Linguistics/philosophy of language
  - Logic
  - Psychology/cognitive science
  - Philosophy of science
  - Artificial intelligence

Linguistics/philosophy of language

- The working assumption in natural language semantics is that the meaning of a sentence lies in its truth-conditions.
- These determine the proposition expressed, and play a key role in theories of semantic composition, discourse pragmatics, and propositional attitudes.
- That is, our answers to questions such as:
  - What if the meaning of "probably A"?
  - What happens when someone asserts A in a discourse?
  - What is it to believe that A?

  refer to the truth-conditions of A, or the proposition expressed.

Linguistics/philosophy of language

- The working assumption in natural language semantics is that the meaning of a sentence lies in its truth-conditions.
- These determine the proposition expressed, and play a key role in theories of semantic composition, discourse pragmatics, and propositional attitudes.
- That is, our answers to questions such as:
  - What if the meaning of "probably A"?
  - What happens when someone asserts A in a discourse?
  - What is it to believe that A?

  refer to the truth-conditions of A, or the proposition expressed.
- However, it turns out to be extremely difficult (some claim, impossible) to specify in which conditions a conditional is true.

- Consider for instance a counterfactual such as:

  (3) If Smith had been elected, he would have cut public spending.

- What must the world be like for (3) to be true?
- (3) seems to talk about an unrealized possibility.

- What about non-counterfactual conditionals like (4)?

  (4) If Smith is elected, he will cut public spending.

- Here, there is a traditional answer: the material analysis.
  If A then B is true if A is false or B is true.

- But, as is well-known, this analysis leads to trouble.
  For instance, (a) is predicted to be equivalent to (b).

  a. It is not true that if Smith is elected, he will cut public spending.
  b. Smith will be elected and he won't cut public spending.

- But clearly someone can believe/assert (a) without necessarily believing that Smith will be elected.

- In fact, many scholars have argued that conditionals don't have truth-conditions at all — and don't express propositions.
- If so, this calls for a serious revision of our linguistic theories, including:
  - compositional semantics: to be able to interpret sentences like (5), where a quantifier scopes over a conditional.

    (5) Most wild boars will attack a predator if threatened.
  - pragmatics: to understand when conditionals can be asserted, and what effects their assertion on a conversational context.
  - propositional attitudes: to understand what it is to believe, want, suppose, ... a conditional.

Logic

- Standard logical consequence is monotonic: $A \models C$ implies $A, B \models C$.
- By contrast, conditionals are not monotonic:

(6) a. If kangaroos had no tails, they would topple over;
b. If kangaroos had no tails but used crutches, they would not topple over.

(7) a. If Alice invites Bob for dinner, he will go.
b. If Alice invites Bob for dinner and then cancels, he won't go.

Understanding the logic of conditionals is tightly linked with understanding non-monotonic reasoning, in particular:

- Belief revision
  Models how an agent's beliefs S change in response to new information.

  a. $S +$ Alice invited Bob $\models$ Bob will go
  b. $S +$ Alice invited Bob $+$ Alice canceled $\models$ Bob won't go

- Default logic
  Models defeasible inferences drawn based on what is normal.

  a. Zack is a Kangaroo $\rightsquigarrow$ Zack lives in Australia
  b. Zack is a Kangaroo, Zack is in a zoo $\not\rightsquigarrow$ Zack lives in Australia

Cognition

- Making hypotheses is one of the most common and most important mental processes. It allows us to run "mental simulations".
- This ability is crucial to decision making. E.g., suppose you are offered an apartment to rent. You reason:

  (8) What if I took it? It is far from the center, so it would take me ages to go to work. On the other hand there are no neighbors, so I could practice playing trumpet. The rent is cheap, so I'd save money...

  Then you assess the outcome and compare it with the alternatives.

- Conditionals are used to describe the outcomes of such process:

  (9) If I took that house, it would take me forever to go to work.

- Conditionals are used to describe the outcomes of such process:

  (9) If I took that house, it would take me forever to go to work.

- And to discuss them (such simulations can be wrong!)

  (10) Actually no, there's a bus which runs directly from there. It would take you about 20 minutes.

- Conditionals are used to describe the outcomes of such process:

  (9) If I took that house, it would take me forever to go to work.

- And to discuss them (such simulations can be wrong!)

  (10) Actually no, there's a bus which runs directly from there.
  It would take you about 20 minutes.

- Conditionals gives us empirical access into this vital cognitive process.

The ability to think conditional thoughts is a basic part of our mental equipment. A view of the world would be an idle, ineffectual affair without them. There's not much point in recognising that there's a predator in your path unless you also realise that if you don't change direction pretty quickly you will be eaten.

(Edgington 1995)

Philosophy of science

- What is a law? How is it different from an accidental generalization?

  a. Every human under stress produces adrenaline.
  b. Every human in this room is under 40.

Philosophy of science

- What is a law? How is it different from an accidental generalization?

  a. Every human under stress produces adrenaline.
  b. Every human in this room is under 40.

- Only laws, and not accidental generalizations, support counterfactual conditional claims:

  a. If my uncle was put under stress, he would produce adrenaline.
  b. #If my uncle was put in this room, he would become under 40.

This connection was exploited by various theorists in both directions:

$\Rightarrow$ Give a theory of counterfactuals and use it to explain causation.

A caused B := "had A not occurred, B would not have occurred"

$\Leftarrow$ Analyze causation and use it to explain counterfactuals.

"if A then B" true $\iff$ B causally follows from A + relevant facts

Either way, the study of (counterfactual) conditionals and (causal) laws are tightly linked.

Another link: dispositional properties

Many properties that are part of our understanding of the world seem to be really counterfactual properties:

- x is fragile $\iff$ if x were struck with some force, x would shatter.

Another link: dispositional properties

Many properties that are part of our understanding of the world seem to be really counterfactual properties:

- x is fragile $\iff$ if x were struck with some force, x would shatter.
- x is red $\iff$ if x were exposed to white light, it would reflect only frequencies in the range...
- Things can be red even in the dark!

Another link: dispositional properties

Many properties that are part of our understanding of the world seem to be really counterfactual properties:

- x is fragile $\iff$ if x were struck with some force, x would shatter.
- x is red $\iff$ if x were exposed to white light, it would reflect only frequencies in the range. . .
- Things can be red even in the dark!
- The problem of how to test these properties is the problem of assessing certain counterfactual conditionals.

A satisfactory definition of scientific law, a satisfactory theory of confirmation or of disposition terms [. . .] would solve a large part of the problem of counterfactuals. Conversely, a solution to the problem of counterfactuals would give us the answer to critical questions about law, confirmation, and the meaning of potentiality.

I am not at all contending that the problem of counterfactuals is logically or psychologically the first of these related problems. It makes little difference where we start if we can go ahead. If the study of counterfactuals has up to now failed the pragmatic test, the alternatives approaches are little better off.

(Goodman '47)

Artificial intelligence

- In order to act intelligently, a robot needs to be able to reason about the consequences of its actions:
  - to come up with a strategy to achieve goals;
  - to choose which action to take among several options.
- It needs to be able to reason not only about how the world currently *is*, but also about how the world *would be if certain changes were performed*.
- That is, it needs to be able to carry out ontic conditional reasoning.
- Classical logic alone is not sufficient to this purpose.

An example:

- A light bulb is controlled by two switches, A and B.
  The light is on iff the switches are in the same position.

$$\text{On} \leftrightarrow (\text{A-up} \leftrightarrow \text{B-up})$$

- Currently switch A is up and switch B is down.

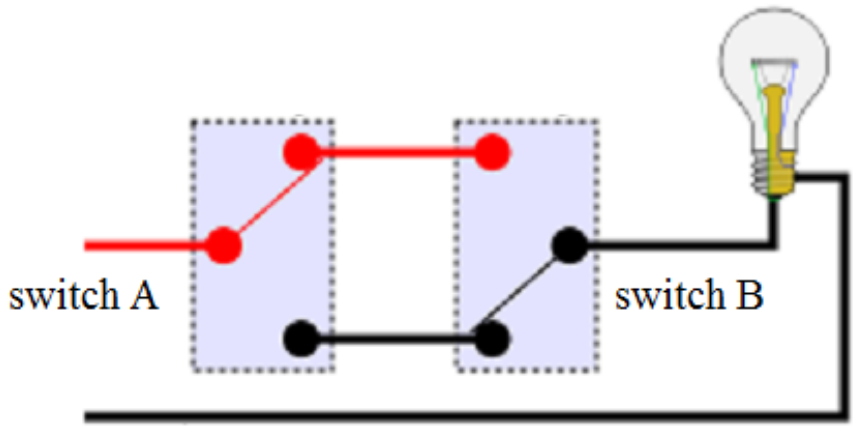

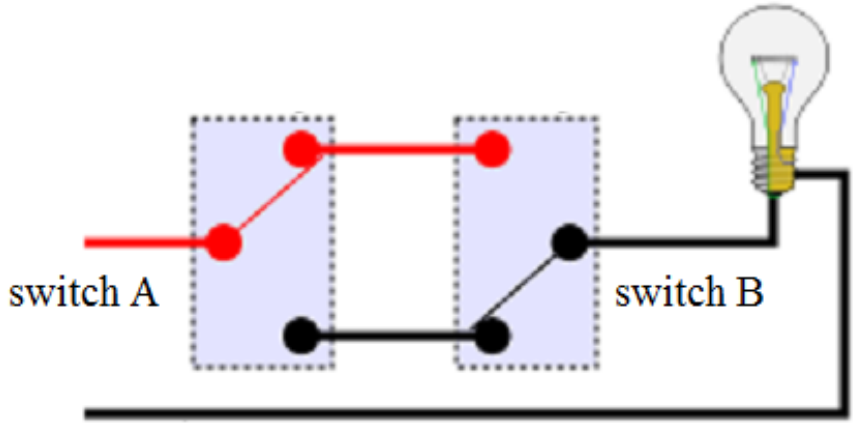

- Now consider a robot whose task is to turn on the light: it should be able to see that this can be achieved, e.g., by toggling switch B.
- I.e. it should infer that if B were switched up, the light would turn on.

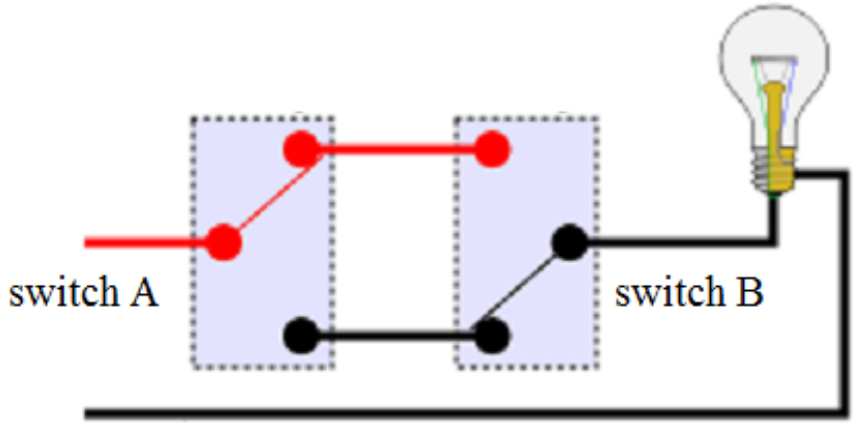

- Now consider a robot whose task is to turn on the light: it should be able to see that this can be achieved, e.g., by toggling switch B.
- I.e. it should infer that if B were switched up, the light would turn on.
- But this doesn't follow by classical logic alone.
- After all, the law is compatible with the following alternative conclusion: if B were switched up, the light would stay off and A would switch down.
- Making the right predictions about ontic conditionals is both crucial to AI, and a non-trivial logical task!

How many conditionals?

Conditionals come in many varieties, differing by mood (indicative vs. subjunctive) and tense.

- Indicatives, e.g.:

a.	If John gets invited, he will go.	(future-oriented)
b.	If John was invited, he went.	(past-oriented)

- Subjunctive:

a.	If John were invited, he would go.	(future-oriented)
b.	If John had been invited, he would have gone.	(past-oriented)

Tense indicates the temporal positions of the events.
But what is the role of the indicative/subjunctive distinction?

A natural idea (e.g., Stalnaker '68):

- indicative and subjunctive conditionals have the same truth-conditions;
- mood is used to express an attitude towards the antecedent, to indicate whether it is being treated as possible (indicative) or not (subjunctive).

a. If John went to Paris, he won't be back until Monday.
$\rightsquigarrow$ he might have gone

b. If John had gone to Paris, he wouldn't be back until Monday.
$\rightsquigarrow$ he didn't go

- Unfortunately, this is not the case, as illustrated by this example from Adams 1970:

  Suppose that on a given occasion three persons, A, B and V, are isolated in a room which is sealed off from the outside. During this time, the third person, V ('the victim'), is murdered by being shot.

  The circumstances are such that only A and B could have done the shooting, though both deny it and accuse the other, and no one else witnessed the murder. An investigation is therefore instituted.

  It establishes that A had in fact a very strong motive for wanting V dead, and furthermore had a gun which might well have been the murder weapon. It also shows that B had no known motive to murder V, and it fails to that B might have had a gun at the time. On this basis it is concluded that A was the murderer.

Given that B appear to have had neither motive nor means to shoot V, it seems preposterous to say:

(11) If A hadn't shot V, B would have.

On the other hand, since only A and B were in the room with V, the following seems true:

(12) If A didn't shoot V, then B did.

Given that B appear to have had neither motive nor means to shoot V, it seems preposterous to say:

(11) If A hadn't shot V, B would have.

On the other hand, since only A and B were in the room with V, the following seems true:

(12) If A didn't shoot V, then B did.

Similarly, we judge the following differently:

(13) a. If A hadn't shot V, then V would still be alive. *True*
b. If A didn't shoot V, then V is alive. *False*

A concrete version of the example. Context:

- On November 22nd, 1963, John Kennedy was shot in Dallas.
- The governmental commission concluded that Lee Oswald he was the murderer, and that he acted alone, not as part of a conspiracy.
- Suppose we are inclined to trust this conclusion.

A concrete version of the example. Context:

- On November 22nd, 1963, John Kennedy was shot in Dallas.
- The governmental commission concluded that Lee Oswald he was the murderer, and that he acted alone, not as part of a conspiracy.
- Suppose we are inclined to trust this conclusion. Consider:

(14) If Oswald hadn't shot Kennedy, someone else would have.

A concrete version of the example. Context:

- On November 22nd, 1963, John Kennedy was shot in Dallas.
- The governmental commission concluded that Lee Oswald he was the murderer, and that he acted alone, not as part of a conspiracy.
- Suppose we are inclined to trust this conclusion. Consider:

(14) If Oswald hadn't shot Kennedy, someone else would have. F

A concrete version of the example. Context:

- On November 22nd, 1963, John Kennedy was shot in Dallas.
- The governmental commission concluded that Lee Oswald he was the murderer, and that he acted alone, not as part of a conspiracy.
- Suppose we are inclined to trust this conclusion. Consider:

(14) If Oswald hadn't shot Kennedy, someone else would have. F

(15) If Oswald didn't shoot Kennedy, someone else did.

A concrete version of the example. Context:

- On November 22nd, 1963, John Kennedy was shot in Dallas.
- The governmental commission concluded that Lee Oswald he was the murderer, and that he acted alone, not as part of a conspiracy.
- Suppose we are inclined to trust this conclusion. Consider:

(14)	If Oswald hadn't shot Kennedy, someone else would have.	F
(15)	If Oswald didn't shoot Kennedy, someone else did.	T

Conclusion:
indicative and subjunctive conditionals do not just signal different attitudes, but can come apart in truth-value, and therefore express different propositions.

A concrete version of the example. Context:

- On November 22nd, 1963, John Kennedy was shot in Dallas.
- The governmental commission concluded that Lee Oswald he was the murderer, and that he acted alone, not as part of a conspiracy.
- Suppose we are inclined to trust this conclusion. Consider:

(14)	If Oswald hadn't shot Kennedy, someone else would have.	F
(15)	If Oswald didn't shoot Kennedy, someone else did.	T

Conclusion:
indicative and subjunctive conditionals do not just signal different attitudes, but can come apart in truth-value, and therefore express different propositions.

Let us examine more carefully why this difference arises.

The ontic interpretation

- What if Oswald had not shot Kennedy?
- In supposing this, we are keeping our actual beliefs fixed (e.g., that Oswald did kill Kennedy, and acted alone) and manipulating the world.
- We are imagining that things had been different than they are in some ways, and drawing conclusions on that basis.
- Since we think Oswald was the only shooter, in this alternative scenario where Oswald doesn't shoot, we conclude that no killing takes place.

The epistemic interpretation

- What if Oswald did not shoot Kennedy?
- In supposing this, we are holding the world fixed (we are still talking about the actual world) and manipulating our beliefs about it.
- We are supposing to believe that Oswald in fact didn't shoot (the official conclusion was wrong!) and trying to square this with the rest of our beliefs.
- Since we have very strong evidence that the shooting in fact took place, we conclude that someone else must have shot.

Summing up

- "If it were A, it would be B" $\approx$
  if the world were one where A is true (which it is not), B would be true.
- "If it is A, then it is B" $\approx$
  if the world is in fact one where A is true (which it might be), B is true.

- Does the (syntactic) indicative/subjunctive distinction reliably mark the (semantic) ontic/epistemic distinction?
- In other words, is it generally true that that:
  - indicative conditionals express the result of epistemic hypotheses;
  - subjunctive conditionals express the result of ontic hypotheses.

- Does the (syntactic) indicative/subjunctive distinction reliably mark the (semantic) ontic/epistemic distinction?
- In other words, is it generally true that that:
  - indicative conditionals express the result of epistemic hypotheses;
  - subjunctive conditionals express the result of ontic hypotheses.
- This is a debated topic in the literature.
- I think that this is plausible for past-oriented conditionals, but that future-oriented indicative conditionals can have both readings.

  (16) If Alice goes to the party, Bob will not go.

- The difference between the two interpretations becomes hard to detect when assumptions about the world are not counterfactuals.

Future-oriented indicatives can express ontic conditionals

Consider this example by Edgington:

> *"Don't go in there", I say, "If you go in you will get hurt". You look sceptical but stay outside, when there is large crash as the roof collapses. "You see", I say, "if you had gone in you would have got hurt. I told you so."*

Future-oriented indicatives can express ontic conditionals

Consider this example by Edgington:

> *"Don't go in there", I say, "If you go in you will get hurt". You look sceptical but stay outside, when there is large crash as the roof collapses. "You see", I say, "if you had gone in you would have got hurt. I told you so."*

- Here, the meaning of the initial indicative conditional is later re-cast, accurately, by means of a subjunctive past conditional.
- The reading of "if you had gone in you would have got hurt" is ontic: we are imagining a world in which things went differently than they did.
- The original conditional "if you go in you will get hurt" must thus share this ontic reading. This is plausible, as there is a causal, and not merely epistemic, connection between going in and getting hurt.

Future-oriented indicatives can express epistemic conditionals

Another example, also inspired by Edgington:

> *There are two prisoners, Smith and Jones, and we have reliable information that one of them will try to escape tonight. Jones is a docile and unadventurous chap; Smith is just the opposite.*

I say:

(17) a. It is probably Smith who will try to escape.

b. But if Smith doesn't try to escape, Jones will.

Future-oriented indicatives can express epistemic conditionals

Another example, also inspired by Edgington:

> *There are two prisoners, Smith and Jones, and we have reliable information that one of them will try to escape tonight. Jones is a docile and unadventurous chap; Smith is just the opposite.*

I say:

(17) a. It is probably Smith who will try to escape.
b. But if Smith doesn't try to escape, Jones will.

As a matter of fact, Smith does try to escape. Can I now assert:

(18) If Smith had not tried to escape, Jones would have.

Future-oriented indicatives can express epistemic conditionals

Another example, also inspired by Edgington:

> *There are two prisoners, Smith and Jones, and we have reliable information that one of them will try to escape tonight. Jones is a docile and unadventurous chap; Smith is just the opposite.*

I say:

(17) a. It is probably Smith who will try to escape.
b. But if Smith doesn't try to escape, Jones will.

As a matter of fact, Smith does try to escape. Can I now assert:

(18) If Smith had not tried to escape, Jones would have.

I think not. If Smith had not tried to escape, it would have been a quiet night.

The reading of (17-b) is clearly epistemic: the claim is not based on a causal connection between Smith's non-escape and Jones's escape. Accordingly, it is not legitimate to cast (17-b) as a subjunctive.

Delimiting our scope

- Ontic and epistemic readings of conditionals pose different theoretical challenges; in both cases, the literature is enormous.
- In this course, we will focus on ontic readings of conditionals.
- If time permits, we will discuss some challenges arising from epistemic conditionals at the end of the course.

- In the literature, it is common to drawn the line between subjunctive and indicative conditionals, with the former referred to as “counterfactuals”.
- I prefer the semantic distinction: as I’ve argued, some indicative conditionals involve the same kind of conditional reasoning as counterfactuals do, and they require the same kind of account.

- In the literature, it is common to drawn the line between subjunctive and indicative conditionals, with the former referred to as "counterfactuals".

- I prefer the semantic distinction: as I've argued, some indicative conditionals involve the same kind of conditional reasoning as counterfactuals do, and they require the same kind of account.

- In addition, the common term "counterfactuals" is a bit of a misnomer. Although such conditionals are typically used when the antecedent is known to be false, this is not always the case:

  (19) If Jones had taken arsenic, he would have shown just exactly those symptoms which he does in fact show. (Anderson, 1951)