

PHILOSOPHY OF LOGIC AND LANGUAGE

WEEK 7: LOGICAL CONSEQUENCE

JONNY MCINTOSH

INTRODUCTION

In the past couple of weeks, we've been looking at Tarski's work on truth and, relatedly, work on the Liar Paradox.

This week, and next, we'll look at work by Tarski and others on **LOGICAL CONSEQUENCE** and **LOGICAL CONSTANTS**.

(Owing to the strike, there was no lecture last week, and will probably be no lecture next week. But I'll put slides online.)

TARSKI'S ACCOUNT

The truth definitions we've been looking provide definitions of truth for **INTERPRETED** languages, whose sentences have meanings that determine truth-values.

These include languages such as the **LANGUAGE OF ARITHMETIC**. In these, **NON-LOGICAL** expressions such as '0', 'S' ('the successor of'), '+', and 'x' have fixed meanings.

In later work, Tarski showed how we can also provide definitions of truth in a **MODEL** for **UNINTERPRETED** languages, whose sentences *don't* have meanings that determine truth-values.

These include languages such as L_1 , L_2 , and L_3 of first year. In these, non-logical expressions such as 'P', 'Q', 'a', and 'b' do *not* have fixed meanings.

Roughly, a model for a language specifies just enough information about its non-logical vocabulary for assigning truth values to each of the sentences of the language.

A bit more precisely, a **MODEL** for a language L is a nonempty domain D plus an appropriate assignment of denotations from D to the basic non-logical expressions of L .

For example, constants (names) might be assigned objects in D and n -place predicates might be assigned sets of n -tuples of objects in D .

We can then define truth in a model for an uninterpreted language by abstracting from definitions of truth (*simpliciter*) that we give for interpreted languages with the same vocabulary.

In the case of the uninterpreted language of predicate logic, the result is the definition of truth in a model (or **STRUCTURE**) that you're familiar with from 1st year.

And using this, we can go on to define the notion of **LOGICAL CONSEQUENCE** (and related notions such as **LOGICAL VALIDITY** and **LOGICAL TRUTH**).

A sentence ϕ is a **LOGICAL CONSEQUENCE** of a set Γ of sentences IFF ϕ is true in every model in which every member of Γ is true.

To better appreciate the merits of this, let's think a little bit about what we *want* from a definition of the notions of logical consequence.

LOGICAL CONSEQUENCE

What is it for a conclusion, ϕ , to be a logical consequence of a set of premises, Γ ?

ARGUMENT 1

1. Everyone smokes and everyone drinks
2. Everyone smokes and drinks

The premise of this argument might not be true, but one thing we seem to be sure of is that, *if* it is true, the conclusion is also true.

Otherwise put: it is not the case that the premises are all true and the conclusion is false. We'll say that such an argument is **TRUTH PRESERVING**.

In order for the conclusion of an argument to be a logical consequence of the premises, it is *necessary* that the argument be truth preserving. But it is obviously not *sufficient*.

ARGUMENT 2

1. London is the capital of the U.K.
2. So Paris is the capital of France

So what else is needed? There are broadly speaking two ideas. One appeals to the notion of **NECESSITY**. The other appeals to the notion of **FORMALITY**.

NECESSITY

The first thought: the conclusion of an argument is a logical consequence of its premises IFF the argument is, in some sense, **NECESSARILY** truth preserving.

That is to say, is in some sense not **POSSIBLE** for the premises to be true and the conclusion false.

But: the conclusion of an argument is a logical consequence of its premises IFF the argument is, in *what* sense, necessarily truth preserving?

One idea: the conclusion of an argument is a logical consequence of its premises IFF there is no **POSSIBLE WORLD** in which the premises are true and the conclusion is false.

This marks a difference between arguments 1 and 2. Although both are truth preserving, only argument 1 is, in this sense, necessarily truth preserving.

But it does not mark a difference between argument 1 and other arguments where, intuitively, the conclusion is *not* a logical consequence of the premises.

ARGUMENT 3

1. This cup contains water.
2. This cup contains H₂O.

There is no possible world in which the premise of argument 3 is true and its conclusion is false, but its conclusion is not a *logical* consequence of its premises.

Notice that, while 'water' and 'H₂O' necessarily refer to the same substance, it is not part of their *meanings* that they refer to the same substance.

So perhaps: the conclusion of an argument is a logical consequence of its premises IFF it is not **CONCEPTUALLY** possible for the premises to be true and the conclusion false.

Yet the relevant notion of *conceptual* possibility is murky. It depends on the analytic-synthetic distinction, famously attacked by Quine.

An alternative: the conclusion of an argument is a logical consequence of its premises IFF it is **KNOWABLE A PRIORI** that the argument is truth-preserving.

But this raises a host of issues. What is *a priori* knowledge? Do we have any? And if so, how is it even possible for us to have it?

Insofar as an account of how it is possible to have to *a priori* knowledge depends on the analytic/synthetic distinction, it's unclear this fares any better.

Moreover, while appeals to conceptual necessity and *a priori* knowability may help to distinguish argument 1 from arguments 2 and 3, they don't help with...

ARGUMENT 4

1. John is a bachelor.
2. So John is not married.

There is presumably no conceptual possibility in which the premise of argument 4 is true and its conclusion is false.

Moreover, that the conclusion of argument 4 is true if its premise is true seems to be something that we can know *a priori*.

Yet the conclusion of argument 4 does not seem to be a *logical* consequence of its premise.

In sum: appeal to the notion of necessity doesn't seem to get us any closer to *sufficient* conditions on a conclusion's being a logical consequence of a set of premises.

FORMALITY

A different thought: while arguments like **ARGUMENT 4** are truth-preserving, they are truth-preserving, not in virtue of their *form*, but rather in virtue of their *matter*.

To see the thought here, notice that **ARGUMENT 4** is an instance of a certain *pattern* of argument, obtained by replacing its non-logical expressions with schematic letters:

1. a is an F
2. So, a is not a G

And other instances of the same pattern are *not* truth-preserving:

- ARGUMENT 5**
1. Theresa is an MP
 2. So, Theresa is not a Conservative

By contrast, **ARGUMENT 1** is an instance of a different pattern of argument:

1. Every F is a G and every F is an H
2. So, every F is a G and an H

And intuitively, other instances of *this* pattern are truth-preserving.

The idea, then, is that if the conclusion of an argument is a logical consequence of its premises, the argument is truth-preserving in virtue of its **LOGICAL FORM**, where...

...the **LOGICAL FORM** of an argument (or sentence) is the pattern of argument (or sentence) obtained by replacing its non-logical expressions with schematic letters.

When an argument (or sentence) is an instance of a certain logical form, we may say that it is a **SUBSTITUTION INSTANCE** of that form.

In these terms, it seems that the conclusion of an argument is a logical consequence of its premises *only if* every substitution instance of that argument is truth-preserving.

That is to say, it is a *necessary* condition on the conclusion's being a logical consequence of the premises that every substitution instance be truth-preserving.

(This will be accepted by anyone who accepts that there is such a thing as the logical form of a sentence, and so of an argument.)

Can we say something stronger? Can we say that it is also a *sufficient* condition on the conclusion's being a logical consequence of the premises?

This is the **SUBSTITUTIONAL** conception of logical consequence:

- The conclusion of an argument is a logical consequence of its premises *if and only if* every substitution instance of that argument is truth-preserving.

A worry: it may be that every substitution instance of an argument is truth-preserving because of expressive limitations of the language in which it is formulated.

For example, in a language that contains just one name, a , which denotes the number 2, and one predicate, F , which denotes even numbers, the sentence Fa will count as a logical truth.

Another worry: it may be that every substitution instance of an argument is truth-preserving because of contingent facts about the cardinality of the universe.

For example, since there are more than two objects, the sentence ' $\exists x \exists y x \neq y$ ', which contains no non-logical expressions, will count as a logical truth too.

Tarski's account of logical consequence can be understood to belong to the same tradition as the substitutional conception, but it is slightly different.

Both explain logical consequence in terms of **THE ABSENCE OF COUNTER-EXAMPLES**. But they offer different accounts of the *range* of potential counter-examples.

On the substitutional conception, a counter-example is a substitution instance of an argument's logical form whose premises are all true and whose conclusion is false.

For Tarski, a counter-example is rather a model (or structure) in which the premises of the argument are all true and the conclusion is false.

Since a model pairs non-logical expressions, not with other expressions in the language, but rather with appropriate denotations from the domain, this addresses the first worry.

The translation of 'Two is even', for example, will turn out to be false in some models that pair the translation of 'two' with the number 3.

And since different models have different domains of quantification, with different cardinalities, it also addresses the second worry.

Since there are domains with just one object, for example, there are models in which the sentence ' $\exists x \exists y x \neq y$ ' comes out as false.

PROBLEMS

Tarski's model-theoretic account of logical consequence is appealing, then. But it also faces a number of problems.

First, it relies on a distinction between the logical and non-logical expressions of a language. (A model pairs *non-logical* expressions with appropriate values.)

But *how*, exactly, are logical expressions or constants to be distinguished from non-logical ones?

This is **THE PROBLEM OF LOGICAL CONSTANTS**. It's the topic of next week's lecture, so I won't say any more about it today.

John Etchemendy famously offers two objections designed to show that Tarski's account of logical consequence is theoretically inadequate.

CONCEPTUAL INADEQUACY

Etchemendy's first objection is that the model-theoretic account of logical consequence is **CONCEPTUALLY** inadequate.

On the model-theoretic account, an argument is logically valid IFF there are no models in which its premises are true and its conclusion is false.

According to Etchemendy, this leaves something essential out of account: the logical validity of an argument provides a **GUARANTEE** that the argument is truth-preserving.

It perhaps *follows* from the fact that an argument is logically valid that there are no models in which its premises are true and its conclusion is false.

But its logical validity does not *consist* in there being no models in which its premises are true and its conclusion is false.

According to Etchemendy, the model-theoretic account of logical consequence thus makes a mistake akin to that of mistaking the symptoms of a disease for the disease itself.

EXTENSIONAL INADEQUACY

Etchemendy's second objection is that the model-theoretic account of logical consequence is **EXTENSIONALLY** inadequate.

He thinks the model-theoretic account both **OVERGENERATES**, i.e. declares as logically valid arguments that are *not* logically valid ...

... and that it **UNDERGENERATES**, i.e. declares as logically *invalid* arguments that are not logically invalid.

Etchemendy's focus is on *overgeneration*. But he does not think that the model-theoretic account overgenerates in first-order logic.

Thanks to an argument from George Kreisel (1967), known as the **SQUEEZING ARGUMENT**, it can be shown that the model-theoretic account does *not* overgenerate in first-order logic.

In order to find examples of arguments which are truth-preserving in all models but not logically valid, Etchemendy therefore focuses on second-order logic.

The argument turns on the **CONTINUUM HYPOTHESIS**. This is the hypothesis that there is no set whose cardinality is between that of the integers and the real numbers.

It is possible to use nothing but logical expressions of second-order logic to formulate a sentence which is true in all second-order models IFF the continuum hypothesis is true.

Call this sentence **S**. Its negation, $\neg S$, is true in all second-order models IFF the continuum hypothesis is false.

Now consider the following arguments:

ARGUMENT A

1. Donald Trump is a Republican
2. So, **S**

ARGUMENT B

1. Donald Trump is a Republican
2. So, $\neg S$

If the continuum hypothesis is true, then **S** is true in all models, and **ARGUMENT A** is declared logically valid.

If the continuum hypothesis is false, then $\neg S$ is true in all models, and **ARGUMENT B** is declared logically valid.

Either way, one of the two arguments gets declared logically valid. But, Etchemendy claims, neither of them is in fact logically valid.

Why not? The thought seems to be that they can only be logically valid if either the continuum hypothesis or its negation is a logical truth.

But it is not the case that either the continuum hypothesis or its negation is a logical truth.

SUMMARY

We've looked at the intuitive notion of logical consequence, and seen that it seems to involve the notion of **FORMALITY**.

This is nicely captured by Tarski's model-theoretic account of logical consequence. But Tarski's account faces various problems.

Next week we'll focus on the problem of logical constants, but along the way look at an alternative to the model-theoretic account.