

PY1003 Introduction to Logic
Lecture 3: Sentential Logic Revisited II

A	$\neg A$
T	F
F	T

A negation is true when and only when the *negated* sentence is false.

A	B	$A \wedge B$
T	T	T
T	F	F
F	T	F
F	F	F

A conjunction is true when and only when both *conjuncts* are true.

A	B	$A \vee B$
T	T	T
T	F	T
F	T	T
F	F	F

A disjunction is false when and only when both *disjuncts* are false. (We use an 'inclusive' disjunction, which is true when both disjuncts are true.)

A	B	$A \rightarrow B$
T	T	T
T	F	F
F	T	T
F	F	T

A conditional is false when and only when the *antecedent* is true and the *consequent* false.

(Therefore, when A means '2+2=5' and B means 'The sky is pink', $A \rightarrow B$ is true. This is sometimes called the 'paradox of material implication', because the conditional involved, \rightarrow , is sometimes known as the 'material conditional'. Some people doubt whether the material conditional is a good translation of the English conditional 'If ... then ...' for this sort of reason.)

A	B	A ↔ B
T	T	T
T	F	F
F	T	F
F	F	T

A biconditional is true when and only when its antecedent and consequent have the same truth-value.

Interpretations

An *interpretation* of a Sentential argument is an assignment of truth-values to all the atomic sentences that appear in the argument.

Similarly, an interpretation of a Sentential formula or set of formulae is an assignment of truth-values to all the atomic sentences that appear in the target sentence or sentences.

If you like, you can think of the interpretation of a sentence as a row of its truth-table. The rows of a truth-table are set up to ensure that every possible interpretation is covered in exactly one row.

Truth-tables for more complex compound formulae

Consider the sentence

- (1) Either the president will accept the bribe or else he will refuse; and there will be a coup.

(Cf: Either the president will accept the bribe, or he will refuse and there will be a coup.)

We shall adopt the following translation scheme:

A: The president will accept the bribe

B: There will be a coup

We can now translate (1) into the language of Sentential logic:

(1') $(A \vee \neg A) \wedge B$

The truth-table for (1') is as follows:

A	B	(A	v	¬	A)	∧	B
T	T	t	t	f	t	T	t
T	F	t	t	f	t	F	f
F	T	f	t	t	f	T	t
F	F	f	t	t	f	F	f
[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]

The order in which we construct this truth-table is very important. We begin by filling in the columns for the atoms A and B, i.e. columns 1 and 2, making sure we cover every possible interpretation. It is then easy to fill in columns 3, 6, and 8 by referring back to 1 and 2.

When we come to working out the truth value of the compound sentences, it is crucial to start with the connective which has the *smallest scope*.

Thus we fill in the remaining columns in the following order: 5, 4, and 7. The final column that we fill in, 7, corresponds to the connective which has the widest scope (the main connective). The truth-values that we write under this connective correspond to the truth-values of the whole formula. We use capital letters here to indicate this.

Comparing column 7 with columns 1 and 2, we can read off that our original sentence (1) is true if and only if it is true that there will be a coup.

Some more truth-functional connectives

So far, we have mainly considered the English sentential connectives “it is not the case that...”, “or”, “and”, “if... then...” and “if and only if”, which we represent with the symbols \neg , \vee , \wedge , \rightarrow and \leftrightarrow . But there are many other sentential connectives in English:

- The president will be assassinated *unless* he bribes the leader of the opposition.
- *It is neither the case that* the president will be assassinated *nor* that there will be a coup.
- The coup will succeed *but* the president won't be assassinated.
- The coup will succeed *only if* the president is assassinated.
- *Only if* Lister washes his shirt will Kochanski be impressed.

We could introduce new symbols to represent these extra sentential connectives. But this would be unnecessary. For each of these connectives can be adequately expressed using the symbols that we already have. The translations of the sentential connectives above are as follows:

A unless B	$\neg B \rightarrow A$
It is neither the case that A nor the case that B	$\neg (A \vee B)$ or $\neg A \wedge \neg B$
A but B	$A \wedge B$
A only if B	$A \rightarrow B$
Only if A, B	$B \rightarrow A$

In fact, we don't even need all the symbols that we have in Sentential. For instance, instead of introducing a special symbol for the biconditional, we could represent the claim 'A if and only if B' as follows:

$$(A \rightarrow B) \wedge (B \rightarrow A)$$

In order to check that this is a good translation we can compare the truth table for it with the truth table for

$A \leftrightarrow B$. This would show us that $(A \rightarrow B) \wedge (B \rightarrow A)$ is true in exactly the same circumstances that $A \leftrightarrow B$ is true. Which is to say that the two formulae are *logically equivalent*.

Two formulae are *logically equivalent* iff they have the same truth-value on all interpretations, i.e. on every assignment of truth-values to the atomic sentences.

Trivially, every formula is logically equivalent to itself.

Many English sentences can be translated in various logically equivalent ways. Logically equivalent formulae (or *logical equivalents*) are often just as good as translations of the same English sentence, **except** where one is much better for bringing out the logical form of the corresponding English sentence.

For instance:

Neither Marcus nor Philip went to the party
can be just as well translated in either of the following logically equivalent ways:

$$\neg A \wedge \neg B$$

$$\neg(A \vee B)$$

But recall our earlier example:

(1) Either the president will accept the bribe or else he will refuse; and
there will be a coup

which we translated as $(A \vee \neg A) \wedge B$.

This turned out to be true in exactly the same circumstances as:

There will be a coup
which we translated as B.

So it turned out that $(A \vee \neg A) \wedge B$ is logically equivalent to B . Still, B would be a bad translation of (1) because it does not make transparent the logical form of (1).

As Sentential developed, it was decided that the five symbols we're using represent a good trade-off between clarity of logical form and ease of expression. That's why we use these five, and not more or fewer.

Many English connectives are not truth-functional

Consider some more English connectives:

- *The president believes that* there will be a coup
- *It is possible that* the president will be shot
- The president was shot *after* he accepted the bribe.

These sentential connectives are not truth-functional.

As a result, they cannot be expressed using the language of Sentential logic except as atomic sentences. For example, consider the compound sentence "The president was shot after he accepted the bribe." If we write "A" for "the president was shot" and B for "he accepted the bribe", we can express this sentence in Loglish as follows: "A after B". When is this sentence true? Well, clearly it is true only when both A and B are true. If the president was not shot, or he did not accept the bribe, it would be false that the president was shot after he accepted the bribe. But this is not enough. Suppose that the president was shot *before* he accepted the bribe. Then it would be false that "A after B" even though both A and B are true.

In other words, the connective *after* is not truth-functional. The truth-value of compounds involving it does not depend entirely on the truth-values of the atomic sentences it governs, but also upon the times that the events described take place. In this module, we will not be dealing with non-truth-functional connectives, although you will encounter some in the more advanced logic courses available in St Andrews.

Truth-tables, validity and counterexamples

1. Ross in his office or Ross is in the library
2. It is not the case that Ross is in his office
3. Therefore, Ross is in the library

We can determine the validity of this argument using a truth table. We adopt the following translation scheme:

A: Ross is in his office
 B: Ross is in the library

This enables us to write out the argument in sequent form:

$$A \vee B, \neg A \vdash B$$

Recall that an argument is valid iff it is impossible for all the premises to be true while the conclusion is false. So to assess the validity of this argument, we need to work out whether there are any possible situations in which $A \vee B$ is true and $\neg A$ is also true, but B is false.

A	B	A	v	B	¬	A	B	
T	T	t	T	t	F	t	T	[1]
T	F	t	T	f	F	t	F	[2]
F	T	f	T	t	T	f	T	[3]
F	F	f	F	f	T	f	F	[4]
[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]	

(NB It isn't necessary to number your truth tables – I'm just doing this to demonstrate what's happening here.)

Columns 1 and 2 represent all the ways that the world could be with regard to the truth-values of A and B – that is, they represent all the possible interpretations.

Columns 4, 6, and 8 are the important ones, because they tell us the truth-values of the premises and the conclusion on each interpretation. We can see from the truth-table that there is only one possible situation in which both premises are true. This occurs when A is false and B is true and is represented by row 3 of the table. Looking at across to column 8, we see that in these circumstances the conclusion, too, is true. So we can conclude that the argument is valid. That is, in all possible circumstances in which the premises are true, the conclusion is also true.

Now consider another argument:

1. Ross is either in his office or in the library
2. Ross is in his office
3. Therefore, Ross is not in the library

Using the same translation scheme, we can write this out as follows:

$$A \vee B, A \vdash \neg B$$

And here is the truth table for this argument:

A	B	A	v	B	A	¬	B	
T	T	t	T	t	T	F	t	[1]
T	F	t	T	f	T	T	f	[2]
F	T	f	T	t	F	F	t	[3]
F	F	f	F	f	F	T	f	[4]
[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]	

Columns 4, 6, and 7 are the important ones here, since they show the truth-values of the premises and the conclusion. Looking at row [1], we see that there are circumstances in which the two premises are true, and yet the conclusion is false. This happens when both A and B is true (perhaps Ross' office is in the library!). So we know that the argument is invalid.

If there is an interpretation on which the premises of an argument are true and the conclusion is false, we say that this interpretation is a *counterexample* to the argument.

Tip: once you've found a counterexample, there is no need to complete the rest of the table.

We can *exhibit* (display) the counterexample to the argument above as follows:

I(A) = T

I(B) = T

This means that the required interpretation, which we can call 'I', assigns the truth-value T to both A and B.

An argument is *valid* iff there is no interpretation in which the premises are true and the conclusion is false.

A *counterexample* to an argument is an interpretation which makes the premises true and the conclusion false.

Tautology, contradiction and consistency

A sentence is a *tautology* (or a *tautologous* sentence) iff it is true on all interpretations.

A sentence is a *contradiction* (or a *contradictory* sentence, or an *inconsistent* sentence) iff it is true on no interpretation.

A sentence is *consistent* iff it is true in at least one interpretation.

Note that all tautologies are consistent.

We can use truth-tables to test for tautology, contradiction and consistency. If a sentence is a tautology then the column under the main connective will contain nothing but 'T's.

For instance, the sentence ' $A \vee \neg A$ ' is a tautology:

A	A	\vee	\neg	A
T	t	T	f	t
F	f	T	t	f

If a sentence is contradictory, then the column under the main connective will contain nothing but 'F's. For instance, the sentence ' $A \wedge \neg A$ ' is a contradiction:

A	A	\wedge	\neg	A
T	t	F	f	t
F	f	F	t	f

And if a sentence is consistent then there will be at least one 'T' under the main connective. For instance, the sentence ' $A \rightarrow (A \wedge B)$ ' is consistent:

A	B	A	\rightarrow	(A	\wedge	B)
T	T	t	T	t	t	t
T	F	t	F	t	f	f
F	T	f	T	f	f	t
F	F	f	T	f	f	f

Testing sets of sentences for consistency

Just as a single sentence can be contradictory (inconsistent with itself), so a set of sentences can be *jointly* inconsistent: incapable of all being true together. This can happen even if every individual sentence in the set is consistent. It will happen automatically if one of the individual sentences is a contradiction.

The standard notation for sets uses the curly brackets: { }

Inside the brackets we list the members of the set, separated by semicolons.

Here is a set of English sentences:

{Lister will not wash his shirt; Kochanski will not be impressed; Koalas can climb trees}

And here is a set of Sentential formulae:

{ $A \wedge B$; $\neg B$ }

We can test sets of sentences for consistency using truth tables. For instance, here is the test for the above set of Sentential formulae:

A	B	(A	∧	B)	¬	B
T	T	t	T	t	F	t
T	F	t	F	f	T	f
F	T	f	F	t	F	t
F	F	f	F	f	T	f

To prove consistency, we need an interpretation (a row of the table) where all the sentences in the set are true. In this case, there is no such row. This table proves the inconsistency of the set.

Further reading

At this point, you should be feeling confident about how to construct truth-tables. If you are not, I strongly suggest some extra reading. The following texts are helpful:

Colin Howson, *Logic With Trees*, ch.1, sections 3-8

Graeme Forbes, *Modern Logic*, ch. 3, section 1

(If reading the Forbes, watch out for the different symbolism!)

You should also be confident in using truth-tables to test the validity of an argument, and to determine whether a sentence is tautologous, whether it is contradictory, and whether it is consistent. You should understand the notion of an interpretation of a formula of sentential logic and know how to exhibit counterexamples to invalid arguments.

If you want to do some further reading on this, I'd recommend: Graeme Forbes, *Modern Logic*, ch.3, sections 2 and 3.