Chapter 4: The Logic of Boolean Connectives

§ 4.1 Tautologies and logical truth

Logical truth

We already have the notion of **logical consequence**. A sentence is a logical consequence of a set of sentences if it is impossible for that sentence to be false when all the sentences in the set are true. We will define logical truth in terms of logical consequence.

Suppose a given sentence is a logical consequence of **every** set of sentences. That means that it is impossible for that sentence to be false – it comes out true in every possible circumstance. Hence:

> A sentence is a **logical truth** if it is a logical consequence of **every** set of sentences.

Tautology

A tautology is a logical truth that owes its truth entirely to the meanings of the truth-functional connectives it contains, and not at all to the meanings of the atomic sentences it contains.

For example, Cube(a) ∨ ¬Cube(a). No matter what shape *a* is, this sentence comes out true. And it owes its truth entirely to the meanings of *or* and *not*. You could replace Cube with any other predicate and a with any other name, and the resulting sentence would still be true. Indeed, you could replace Cube(a) with any other sentence and the resulting sentence would still be true.

Tautologies and truth tables

To show that an FOL sentence is a tautology, we construct a **truth table**. Look at the example of the table for Cube(a) ∨ ¬Cube(a) on p. 96.

Features of truth tables

The number of rows in the table for a given sentence is a function of the number of atomic sentences it contains. If there are *n* atomic sentences, there are 2^n rows.

Each row represents a possible assignment of truth values to the component atomic sentences.

On each row, the values of the atomic sentences determine the values of the compounds of which they are components. The values of the compounds of atomic sentences in turn determine the values of the larger compounds of which they are components. In the end, a unique value for the entire sentence is determined on each row.

A **tautology** is a sentence that comes out true on every row of its truth table.

Do the **You try it** on p. 100: Open the program Boole and build the truth table. You will confirm that ¬(A ∧ (¬A ∨ (B ∧ C))) ∨ B is a tautology.

Tautologies, logical truths, and Tarski's World necessities

When we looked at the sentence Cube(a) ∨ ¬Cube(a), we noted that it owes its truth entirely to the meanings of *or* and *not*. You could replace Cube (both occurrences, of course) with any other predicate, and the resulting sentence would still be true. Indeed, you could replace the two occurrences of Cube(a) with any other sentence and the resulting sentence would still be true.

Contrast this with Cube(a) ∨ Tet(a) ∨ Dodec(a). Although this sentence always comes out true in Tarski's World, we can imagine a circumstance in which it is not true: suppose that *a* is a sphere. So this sentence is not even logically true. We can say that it is a **Tarski's World necessity**, because it comes out true in every world in Tarski's World. (It is a special feature of Tarski's World that there are no objects other than cubes, tetrahedra, and dodecahedra.)

So Tarski's World necessities form a large set of sentences that includes the tautologies as a (smaller) part: every tautology is a Tarski's World necessity, but not conversely.

Note that there are some necessary truths that are not tautologies, but don't depend for their truth on any special features of Tarski's World. For example:

¬(Larger(a, b) ∧ Larger(b, a))

This is not a tautology, for it depends on the meaning of the predicate *larger than*. But its necessity is not limited to Tarski's World, for it can *never* be true that both *a* is larger than *b* and *b* is larger than *a*.

Why Boole can't identify all logical truths

Boole is sensitive to the meaning of the truth-functional connectives, but not to the meanings of the predicates contained in atomic sentences. (In particular, Boole does not recognize the meaning of the identity symbol =, nor does Boole recognize the meanings of the quantifier symbols ∀ and ∃ that we'll be studying in chapter 9.)

So when Boole sees a sentence like ¬(Larger(a, b) ∧ Larger(b, a)), it does not "see" the predicate *larger than*. Instead, all Boole sees is the negation of a conjunction of two different atomic sentences. In effect, all Boole sees is sentence of the form ¬(P ∧ Q). And when Boole sees this sentence, it thinks, "I know how to make this sentence false—I just assign **T** to P and **T** to Q. That makes P ∧ Q true, and so it makes ¬(P ∧ Q) false." Since Boole can't "see inside" the atomic sentences and doesn't understand the predicates they contain, he doesn't know that it's impossible for both Larger(a, b) and Larger(b, a) to be true.

Now when it comes to tautologies, Boole rules! So any logical truth that Boole doesn't recognize as coming out true in every circumstance is a non-tautology. A row on a truth-table that contains a **T** under the main connective, then, may not represent a genuine logical possibility.

We have discovered that there is a set of logical truths that falls **between** the tautologies and the Tarski's World necessities. It is best to picture the situation in terms of a nested group of concentric circles ("Euler circles") collecting together a subset of all the true sentences:

- The outer, largest, circle: the Tarski's World necessities (sentences that are "TW-necessary"). It also contains the contents of all the inner circles.
- The next largest circle: the logical truths or logical necessities.

- The innermost circle: the tautologies ("TT-necessary").

The relationship is depicted graphically on p. 102 in figure 4.1. You should be able to give examples of each kind of necessary truth.

Note that every tautology is also a logical truth, and every logical truth is also a TW-necessity. But the converse is not true: some logical truths are not tautologies, and some TW-necessities are not logical truths.

Three kinds of possibility

Notice that if we are considering **possibility**, rather than necessity, we have a similar nest of Euler circles. The difference is that the TT-possible sentences—the ones that come out true on at least one row of their truth table—are included in the largest circle, and the TW-possible sentences are in the smallest circle. That is, a sentence may be TT-possible without being logically possible or TW-possible, although all TW-possibilities are also logically possible and TT-possible.

Look at exercise 4.10. We are asked to locate these three classes of sentences in an Euler diagram. To see what the circles look like, open Possibility.pdf (on the Supplementary Exercises page of the course web site).

- The outer, largest, circle: the "TT-possible " sentences. It also contains the contents of all the inner circles.
- The next largest circle: the logically possible sentences.
- The innermost circle: the Tarski's World possibilities ("TW-possible " sentences).

Again, you should be able to give examples of each kind of possibility. You can test your understanding of these different kinds of possibility by completing exercise 4-9 (not assigned for homework).

I'd suggest downloading and printing a copy of Possibility.pdf for your notes.

§ 4.2 Logical and tautological equivalence

Logically equivalent sentences

Sentences that have the same truth value in every possible circumstance are *logically equivalent*.

Tautologically equivalent sentences

Logically equivalent sentences whose equivalence is due to the meanings of the truth functional connectives they contain are *tautologically equivalent*.

Tautological equivalence and truth tables

To see whether a pair of FOL sentences are tautologically equivalent, we construct a *joint* truth table for them. The two sentences are tautologically equivalent if they are assigned the same truth value on every row.

Note that sentences may be logically equivalent without being tautologically equivalent. A good example is given on pp. 107-8:

$$a = b \land \text{Cube}(a) \iff a = b \land \text{Cube}(b)$$

These sentences are logically equivalent—there is no possible circumstance in which they could differ in truth value. But they are **not** tautologically equivalent, as the truth table on p. 108 shows.

Note that this truth table contains two rows (rows 2 and 3) that do not represent "real" logical possibilities, although they do represent "truth table possibilities." Row 2, for example, assigns **T** to a = b, **T** to Cube(a), and **F** to Cube(b). This assignment does not represent a real logical possibility, since it is not possible for *a* to be a cube while *b* is not a cube if *a* and *b* are the same block.

How, then, can this be a truth table possibility? The answer is that, as we saw above, Boole can't "see inside" atomic sentences and doesn't understand the predicates they contain. As far as Boole is concerned, a = b, Cube(a), and Cube(b) are just three different atomic sentences, to which it can assign any values it likes.

§ 4.3 Logical and tautological consequence

Consequence is the core notion

Q is a logical consequence of P if it is impossible for P to be true and Q false. That is, there is no possible circumstance in which P is true and Q is false.

Both logical truth and logical equivalence are special cases of logical consequence:

- A sentence is a logical truth if it is a logical consequence of the empty set of sentences.
- Two sentences are logically equivalent if they are logical consequences of one another.

Tautological consequence and truth tables

Q is a tautological consequence of P if in the joint truth table for the two sentences there is no row on which P is true and Q is false.

The relation between logical and tautological consequence

As with tautological truth (and equivalence) vs. logical truth (and equivalence), tautological consequence is a special case of logical consequence. That is, every tautological consequence is also a logical consequence, but the converse does not hold—in some cases, Q might be a logical consequence of P but not a tautological consequence.

Examples

Cube(b) is a logical consequence of a = b ∧ Cube(a), but not a tautological consequence of it. That's because there's no possible circumstance in which a = b ∧ Cube(a) is true and Cube(b) is false. But the truth table for these sentences does not show this, since (as we saw above) it is allowed to assign **T** to a = b, **T** to Cube(a), and **F** to Cube(b). This is a "truth table possibility" that is not a "real possibility."

The same distinction obtains between logical possibility and TT-possibility. The sentence Cube(a) ∧ Tet(a) is TT-possible, since it takes a **T** in row 1. But that TT-possibility is not a "real" possibility, for that row represents the (impossible) case in which *a* is both a cube and a tetrahedron.

We will look at both of these examples again in a moment.

§ 4.4 Tautological consequence in Fitch

Using Fitch to check for consequence

Truth tables are mechanical—there is an automatic procedure (an "algorithm") that always give you an answer to the question whether a sentence is a tautology or whether a sentence is a tautological consequence of another sentence or set of sentences.

Instead of using Boole to construct a truth table, you can use the program Fitch to check whether one sentence is a tautological consequence of a given set of sentences.

Do the **You try it** on p. 114 to see how to do this.

Taut Con, FO Con, and Ana Con

These are three methods, of increasing strength, that Fitch uses to check for consequence.

- **Taut Con** checks to see whether a sentence is a tautological consequence of some others. It pays attention only to the truth functional connectives. It is the weakest procedure of the three because it only catches tautological consequence, and misses the broader notions of consequence.
- **FO Con** checks to see whether a sentence is a "first-order" consequence of some others. It pays attention not only to the truth functional connectives but also to the identity predicate and to the quantifiers.
- **Ana Con** checks to see whether a sentence is an "analytic" consequence of some others. It pays attention not only to the truth functional connectives, the identity predicate, and the quantifiers, but also to the meanings of most (but not all!) of the predicates in the blocks language. This notion comes the closest of the three to that of (unrestricted) logical truth.

If a sentence is a tautological consequence of some others it is clearly also a first-order consequence and an analytic consequence of those sentences. But the converse does not hold—some first-order consequences are not tautological consequences, and some analytic consequences are not first-order consequences.

Examples

- Cube(a) ∨ Cube(b) is a tautological consequence of Cube(a). This is obvious—there is no assignment of truth-values to these sentences that makes Cube(a) true and Cube(a) ∨ Cube(b) false.
- Cube(b) is a first-order consequence, but not a tautological consequence, of a = b ∧ Cube(a). We can check this out, first in Boole (see file Ch4Ex1.tt), then in Fitch (see file Ch4Ex1.prf).
- SameSize(a, b) is an analytic consequence, but not a first-order consequence (and hence not a tautological consequence), of ¬Larger(a, b) ∧ ¬Larger(b, a). We can check this out in Fitch (see file Ch4Ex3.prf).
- Cube(a) ∧ Tet(a) is FO-possible (and hence TT-possible), but not logically possible. We can use Boole (see file Ch4Ex2.tt) to show that it is TT-possible. Notice how, with a little trickery, we can also use Fitch (see file Ch4Ex2.prf) to show both that it is TT-possible and FO-possible, but not logically possible.

A warning about Ana Con

The **Ana Con** mechanism does not distinguish between logical necessity and TW-necessity. That is, it counts at least some "Tarski World" consequences as analytic consequences along with logical consequences more narrowly conceived. An example will make this clear.

- According to **Ana Con**, Cube(b) is an analytic consequence of ¬Tet(b) ∧ ¬Dodec(b). (Obviously, this is not a first-order consequence, and hence not a tautological consequence either.)

This happens because **Ana Con** pays attention not only to the meanings of some of the predicates, but also to some of the special features of Tarski's World. Since in Tarski's World there are only three shapes of blocks, it follows that there cannot be a *Tarski World* in which an object is neither a tetrahedron nor a cube nor a dodecahedron.

But while that may be true for every Tarski World, it does not hold for every possible world. In general, it does not follow logically, from the fact that *b* is neither a tetrahedron nor a dodecahedron, that *b* is a cube—*b* might be a sphere. So this example does not seem to be a logical necessity, but only something weaker—a TW-necessity.

Ana Con also has some other limitations. It misses certain TW-necessities, namely, those involving the predicates Adjoins and Between, which it does not understand. For example, ¬Large(a) is a TW-consequence of Adjoins(a, b), since it is impossible in a Tarski world for a large block to adjoin another block. But **Ana Con** will not recognize this consequence.

Similarly, **Ana Con** does not understand any predicates that are not in the blocks language. Hence, it will not know that Older(b, a) is a logical consequence of Younger(a, b), since these predicates are not in the blocks language. So you must use **Ana Con** with caution!