Symbolic Logic: Grammar, Semantics, Syntax

Logic aims to give a precise method or recipe for determining what *follows from* a given set of *sentences*. So we need precise definitions of 'sentence' and 'follows from.'

1. Grammar: What's a sentence?

What are the rules that determine whether a string of symbols is a sentence, and when it is not? These rules are known as the *grammar* of a particular language. We have actually been operating with two different but very closely related grammars this term: a grammar for *propositional logic* and a grammar for *first-order logic*.

The grammar for **propositional logic** (thus far) is simple:

1. There are an indefinite number of propositional variables, which we have been symbolizing as $\mathsf{A}, \mathsf{B}, ...$ and $\mathsf{P}, \mathsf{Q}, ...$; each of these is a sentence. $\perp$ is also a sentence.

2. If A, B are sentences, then:

   - $\neg\mathsf{A}$ is a sentence,
   - $\mathsf{A} \wedge \mathsf{B}$ is a sentence, and
   - $\mathsf{A} \vee \mathsf{B}$ is a sentence.

3. Nothing else is a sentence: anything that is neither a member of 1, nor constructable via successive applications of 2 to the members of 1, is *not* a sentence.

The grammar for **first-order** logic (thus far) is more complex.
2 and 3 remain exactly the same as above, but 1 must be replaced by something more complex. First, a list of all predicates (e.g. Tet or $\mathsf{SameShape}$) and proper names (e.g. max, b) of a particular language L must be specified. Then we can define:

1_{FOL}. A series of symbols s is an atomic sentence $=$ s is an n-place predicate followed by an ordered sequence of n proper names.

2. Semantics: Consequence in terms of truth and falsity

Here is the basic semantic notion of 'follows from,' which should be familiar from before:

semantic consequence ($\models$) A sentence C (semantically) follows from premises $P_1, P_2, ...$ = there is no case in which $P_1, P_2, ...$ are all true and C is false.

The obvious next question is 'What's a case?' This notion is, on its own, imprecise. But we have seen two ways of making it precise: a case is (i) an allowable arrangement of the Tarski's world checkerboard, or (ii) a row of a truth-table. Substituting (i) in for 'case' above gives us one precise notion of following-from, namely *Tarski's-World consequence*. Substituting (ii) in for 'case' gives us another precise notion of following-from, namely *Tautological Consequence*. These are different consequence relations, but they both are defined in terms of truth and falsity; that is, they are both *semantic*.

3. Syntax: Consequence in terms of rules

We have introduction and elimination rules for $\wedge, \vee, \neg, =$, and $\bot$. We can define a precise notion of *follows from* using these rules: the basic idea is that C follows from $P_1, P_2, ...$ exactly when there is a formal proof of the conclusion from the premises. This is a *syntactic*, not semantic, notion of consequence, since truth and falsity play no role.

The official definition of syntactic consequence is as follows:

syntactic consequence ($\vdash$) A sentence C is a *syntactic consequence* of sentences $P_1, P_2, ...$ = C can be reached from $P_1, P_2, ...$ using *only* the introduction- and elimination-rules.

In other words, there is a formal proof of C from $P_1, P_2, ...$. (As this definition makes clear, whether C is a syntactic consequence of $P_1, P_2, ...$ depends on which specific rules you have in your system.)

Question: Is it possible for C to be a semantic consequence of $P_1, P_2, ...$, but not a syntactic consequence? What about the converse: can you imagine circumstances in which C is a syntactic consequence of $P_1, P_2, ...$, but not a semantic consequence?

Logical truth (semantics) vs. Theorem (syntax)

The textbook often downplays the difference between semantic and syntactic consequence. This is very clear in the 'Remember Box' on p.174. 'Logical truth' was defined earlier as a sentence that is true in every possible case. But how do we know that the last line of a proof with no premises is true in every possible case? (The usual word for the last line of a proof without any premises is a *theorem*.)