2 The Language *S*

2.1 Introducing *S*

We saw in the previous chapter that deductive validity can be understood as a matter of form. Indeed, a number of other logical properties that interest us are also matters of form. The symbolic languages we develop in this book will allow us to represent in great detail the form of sentences and arguments. Using these languages we are able to explore and reason about logical properties in a highly systematic fashion.

2.1.1 Compound Sentences and Truth-Functional Logic

In the course of this book we will be concerned with two aspects of the logical form of sentences. One, which we set aside for a few chapters, has to do with the internal structure of simple sentences and the use of quantity terms such as ‘all’, ‘some’, and ‘none’. The other, which we take up immediately, has to do with the way *compound sentences* are constructed from *simple sentences*.

Simple Sentence:
A *simple sentence* is a sentence that contains one subject and one predicate.

Compound Sentence:
A *compound sentence* is a sentence that either contains one or more simple sentences and at least one compounding phrase, or contains a compound subject or a compound predicate.

It will help to have a few examples:

(1) The Earth is at the center of the Universe
(2) The Sun revolves around the Earth

(1) and (2) are simple sentences. They each have a single subject and a single predicate. We can use these two simple sentences to create some compound sentences:

(3) *It is not the case that* the Earth is at the center of the Universe
(4) The Earth is at the center of the Universe, *and* the Sun revolves around the Earth
(5) *If* the Sun revolves around the Earth *then* the Earth is at the center of the Universe

(6) *If* the Earth is *not* at the center of the Universe, *then* the Sun does *not* revolve around the Earth

(3)–(6) are compound sentences. The subsentences (or component sentences) are in regular type, while the compounding phrases are emphasized. (3) comes from (1) by prefixing ‘it is not the case that’, while (4) and (5) each use (1) and (2) to make longer, compound sentences. Since ‘it is not the case that’ modifies a single component, we call it a unary compounding phrase, while the other two compounding phrases, ‘and’ and ‘if-then’ are binary, requiring two component sentences. In (6) we have the unary ‘not’ compounder (a simplification of ‘it is not the case that’) used together with the binary ‘if-then’ compounder. As you can see from (6), compound sentences can themselves be further compounded.

Here is a slightly different (and often more natural) way of compounding simple sentences:

(7) The Earth *and* the Sun are at the center of the Universe
(8) The Sun is at the center of the Universe *or* revolves around the Earth

(7) and (8) are also compound sentences. Here, however, compounding is achieved through use of a compound subject, as in (7), or a compound predicate, as in (8). The compounding words are emphasized. Often we will want to expand such sentences so they look more like (3)–(6). Thus:

(7′) The Earth is at the center of the Universe *and* the Sun is at the center of the Universe
(8′) The Sun is at the center of the Universe *or* the Sun revolves around the Earth

Here (7′) says exactly the same thing as (7), but the expansion allows us to see exactly what the simple components are. Similarly for (8′). We will frequently use this sort of expansion to make clear the structure of sentences we encounter.

Certain sentences (and arguments) have the logical properties they do as a result of the way they (or their premises and conclusion) are compounded from simple sentences. As can be seen above, and in the examples below, there are many ways of compounding sentences, but we shall be interested here only in what are called *truth-functional compounds*.

Truth-Functional Compound:

A sentence is a *truth-functional compound* iff[1] the truth value of the compound sentence is completely and uniquely determined by (is a function of) the truth values of the simple component sentences. Otherwise, the compound sentence is *non-truth-functional*.

For instance, (3) and (4) are each truth-functionally compound sentences. (3) turns out to be true, and it is true precisely because its simple component, (1), is false. Had

1. ‘iff’ is short for ‘if and only if’.

(1) been true, (3) would have been false—sentences compounded with 'not' (or 'it is not the case that') have the opposite truth value of their component. (4) is formed with 'and', and its simple components are (1) and (2). (4) happens to be false because at least one (indeed, both) of its components is false. Had both (1) and (2) been true, (4) would have been true as well—sentences compounded with 'and' are true when and only when both components are true. We will consider (5) and (6) also to be truth-functional compounds, though the analysis of the 'if-then' is a bit complicated, and we save it for later (see Section 2.1.5).

Not all compounding phrases generate truth-functional compounds. Sometimes information other than the truth values of the simple components is needed to determine the truth value of the whole compound sentence. For instance:

(9) *Ptolemy believed that* the Earth is at the center of the Universe
(10) *Copernicus doubted that* the Sun revolves around the Earth
(11) *Copernicus argued that* the Sun does *not* revolve around the Earth
(12) *Copernicus believed that Ptolemy believed that* the Sun revolves around the Earth
(13) *Ptolemy believed that* the Earth is at the center of the Universe, *and Copernicus argued that* the Sun does *not* revolve around the Earth
(14) The Sun revolves around the Earth *because* the Earth is at the center of the Universe
(15) The Earth was at the center of the Universe *before* the Sun revolved around the Earth

The 'believes that', 'doubted that', and 'argues that' phrases are each non-truth-functional. We cannot determine the truth values of (9)–(13) simply on the basis of the truth values of the components (1) and (2). Rather, even if we know the truth value of the components (false), we still need further historical information about Ptolemy and Copernicus. The 'because' and 'before' compounders in (14) and (15) are also non-truth-functional. The truth values of the components of such compounds do not give any information as to causal or temporal relationships—again, we need some further scientific or historical information in such cases.

Note that, because (13) is an 'and' sentence, its truth value *is* wholly determined by the truth values of its *immediate* components ((9) and (11)). But those components themselves involve non-truth-functional compounds, hence the truth value of (13) is not wholly determined by its *simple* components, (1) and (2). Hence (13) is non-truth-functional.

Truth-Functional Logic:

> *Truth-functional logic* is the logic of truth-functional combinations of simple sentences. It investigates the properties that arguments, sentences, and sets of sentences have in virtue of their truth-functional structure.[2]

Truth-functional logic will be the subject of this and the next two chapters. We will study the properties that arguments and sentences have in virtue of their truth-functional

2. Truth-functional logic is also frequently called sentential or statement logic. I prefer the former since the truth functions are what generate the interesting logical properties. The name of our language will be '***S***,' to remind us of the basic non-logical unit, statements (statement letters to be precise).

structure. We do not, at this point, delve into the internal logical structure of the simple sentences. Rather, we concern ourselves with simple sentences only insofar as (i) they are bearers of one of the two truth values, and (ii) they can be combined truth-functionally. As a result, in developing our language for truth-functional logic, we will take simple sentences as the basic non-logical units. Our name for this language will be '***S***,' which should remind you that sentences are the basic non-logical unit. We will represent simple sentences with uppercase letters (called *statement letters*), and interpret them as having either the truth value true or the truth value false, but not both. To represent compounding phrases, we will introduce symbols called *truth-functional connectives* as a means of generating compound sentences from simpler components.

Before giving a formal specification of the language ***S***, let's take a less technical look. In the remainder of this section I try to use example statements with fairly obvious truth values—in some cases true and in others false—in order to facilitate basic understanding of the truth functions.

2.1.2 Negation—It is not the case that ...

The simplest way to generate a truth-functionally compound sentence is to negate (or deny) a sentence.

(1) Grass is green
(2) Grass is not green
(3) It is not the case that grass is green

Choosing 'G' to translate the simple sentence 'Grass is green', and using '¬' (the negation symbol or hook) for 'not' or 'it is not the case that', we can symbolize the above three sentences as:

(1′) G
(2′) ¬G
(3′) ¬G

Translating (1) is just a matter of writing the designated statement letter, as in (1′). Sentences (2) and (3) are only minutely more complex. Note that despite the slight difference in their phrasing, (2) and (3) receive identical treatment when translated into ***S*** as (2′) and (3′). Moving in the opposite direction, from ***S*** to English, sentence (3) illustrates the "official" English reading of '¬G'. The somewhat cumbersome phrase 'it is not the case that' is used because it can always be placed at the start of an English sentence to unambiguously negate the whole sentence. Of course, when no ambiguity threatens, we can give '¬G' the more natural reading in (2).

As one would expect, the negation of a sentence has a truth value opposite that of the original sentence. That is, negation is the truth function that maps a truth value onto its opposite. We can communicate this clearly with the *characteristic truth table* for negation:

$\mathbb{P}$	$\neg\mathbb{P}$
T	F
F	T

Here '$\mathbb{P}$' is a placeholder for any statement letter or more complex formula (see Section 2.2.3). The left-most column is a list of all possible assignments of truth values to the relevant component sentences (in this case only one sentence is relevant, so there are only two possibilities). For each row, the resulting truth value of the compound is placed in the column under the hook. The truth value of the compound goes under the connective symbol '$\neg$' to distinguish it from the truth value of the component(s) ('$\mathbb{P}$' in this case).

Let us take (1′) to be true—grass is green. Given the interpretation of the '$\neg$', (2′) and (3′) are false. Negation is pretty simple, but, as we will see throughout the book, it interacts in interesting ways with the other logical operators.

2.1.3 Conjunction—Both ... and - - -

Another fairly straightforward way to truth-functionally compound sentences is with 'and'. Here are three new examples:

(1) Grass is green and the sky is blue
(2) Grass and snow are green
(3) Both grass is green and snow is green

Choosing 'B' for 'The sky is blue', 'S' for 'Snow is green', and using '$\wedge$' (the conjunction symbol or wedge) for conjunction, we symbolize the above sentences as:

(1′) $G \wedge B$
(2′) $G \wedge S$
(3′) $G \wedge S$

Translating (1) is a simple matter of placing the wedge between the two statement letters, as in (1′). (2) does not, on the face of it, contain two complete sentences. It does, however, have a compound subject, and it is treated as a contraction of the longer (3). Hence, again, despite the slight difference in their phrasing, (2) and (3) receive identical translation into ***S*** as (2′) and (3′). Moving from ***S*** to English, (3) illustrates the "official" reading of '$G \wedge S$'. This official phrasing will sometimes be used to avoid confusion when things get more complicated.

As one would expect, a conjunction is true when and only when both component sentences (called *conjuncts*) are true. As long as one or more of the components is false, the whole conjunction is false. Here is the characteristic truth table:

$\mathbb{P}$	$\mathbb{Q}$	$\mathbb{P} \wedge \mathbb{Q}$
T	T	T
T	F	F
F	T	F
F	F	F

Again, the left-most columns display the possible truth value assignments. We now have a compound with two components, so we use '$\mathbb{Q}$' as a second placeholder and there are four combinations of the two truth values. The column under the wedge displays the value of the whole conjunction, given the values of the conjuncts listed in the same row. This all means that (1′) is true, while (2′) and (3′) are false.

Note, finally, that conjunction is commutative. Like addition or multiplication in mathematics, changing the order of the components does not change the meaning. Thus, the following are adequate translations of (1)–(3), respectively, and are equivalent to their respective counterparts in (1′)–(3′):

(1″) $B \wedge G$
(2″) $S \wedge G$
(3″) $S \wedge G$

2.1.4 Disjunction—Either ... or - - -

The only mildly tricky thing about disjunction is remembering that we choose to use *inclusive* disjunction as opposed to *exclusive* disjunction. The difference is in how we treat the case in which both components (called *disjuncts*) are true. With *inclusive* 'or' the compound is true if and only if one *or both* disjuncts is true (hence it is false only when both disjuncts are false). With *exclusive* 'or' the compound is true if and only if exactly one (*but not both*) of the disjuncts is true. Though perhaps the exclusive predominates, English usage of 'or' varies with the context between inclusive and exclusive. For us it is a matter of convenience which we choose, as we can always define the other via the one, with the help of conjunction and negation. Because it makes certain other definitions and operations more simple, we choose inclusive disjunction.

(1) Grass is green or the sky is blue
(2) Grass or snow is green
(3) Either snow is green or clouds are red

Using 'R' for 'Clouds are red', and '$\vee$' (the disjunction symbol or vee) for disjunction we symbolize as follows:

(1′) $G \vee B$
(2′) $G \vee S$
(3′) $S \vee R$

Again, translating (1) into (1′) should be straightforward. (2) has a compound subject, but we treat it as a contraction of 'Either grass is green or snow is green'. Thus (2) goes over into (2′). (3), of course, translates into (3′), and (3) uses the "official" reading of the vee. Also note that we interpret (1′) as true, but, had we chosen exclusive disjunction, we would have to interpret it as false. (2′) is true as well, and (3′) is false, since both disjuncts are false. Here is the characteristic truth table for disjunction: