1. Developing a Precise Language

1.1 Starting with sentences

We begin the study of logic by building a precise logical language. This will allow us to do at least two things: first, to say some things more precisely than we otherwise would be able to do; second, to study reasoning. We will use a natural language—English—as our guide, but our logical language will be far simpler, far weaker, but more rigorous than English.

We must decide where to start. We could pick just about any part of English to try to emulate: names, adjectives, prepositions, general nouns, and so on. But it is traditional, and as we will see, quite handy, to begin with whole sentences. For this reason, the first language we will develop is called "the propositional logic". It is also sometimes called "the sentential logic" or even "the sentential calculus". These all mean the same thing: the logic of sentences. In this propositional logic, the smallest independent parts of the language are sentences (throughout this book, I will assume that sentences and propositions are the same thing in our logic, and I will use the terms "sentence" and "proposition" interchangeably).

There are of course many kinds of sentences. To take examples from our natural language, these include:

> What time is it?
>
> Open the window.
>
> Damn you!
>
> I promise to pay you back.
>
> It rained in Central Park on June 26, 2015.

We could multiply such examples. Sentences in English can be used to ask questions, give commands, curse or insult, form contracts, and express emotions. But, the last example above is of special interest because it aims to describe the world. Such sentences, which are sometimes called "declarative sentences", will be our model sentences for our logical language. We know a declarative sentence when we encounter it because it can be either true or false.

1.2 Precision in sentences

We want our logic of declarative sentences to be precise. But what does this mean? We can help clarify how we might pursue this by looking at sentences in a natural language that are perplexing, apparently because they are not precise. Here are three.

> Tom is kind of tall.
>
> When Karen had a baby, her mother gave her a pen.

> This sentence is false.

We have already observed that an important feature of our declarative sentences is that they can be true or false. We call this the "truth value" of the sentence. These three sentences are perplexing because their truth values are unclear. The first sentence is vague, it is not clear under what conditions it would be true, and under what conditions it would be false. If Tom is six feet tall, is he kind of tall? There is no clear answer. The second sentence is ambiguous. If "pen" means writing implement, and Karen's mother bought a playpen for the baby, then the sentence is false. But until we know what "pen" means in this sentence, we cannot tell if the sentence is true.

The third sentence is strange. Many logicians have spent many years studying this sentence, which is traditionally called "the Liar". It is related to an old paradox about a Cretan who said, "All Cretans are liars". The strange thing about the Liar is that its truth value seems to explode. If it is true, then it is false. If it is false, then it is true. Some philosophers think this sentence is, therefore, neither true nor false; some philosophers think it is both true and false. In either case, it is confusing. How could a sentence that looks like a declarative sentence have both or no truth value?

Since ancient times, philosophers have believed that we will deceive ourselves, and come to believe untruths, if we do not accept a principle sometimes called "bivalence", or a related principle called "the principle of non-contradiction". Bivalence is the view that there are only two truth values (true and false) and that they exclude each other. The principle of non-contradiction states that you have made a mistake if you both assert and deny a claim. One or the other of these principles seems to be violated by the Liar.

We can take these observations for our guide: we want our language to have no vagueness and no ambiguity. In our propositional logic, this means we want it to be the case that each sentence is either true or false. It will not be kind of true, or partially true, or true from one perspective and not true from another. We also want to avoid things like the Liar. We do not need to agree on whether the Liar is both true and false, or neither true nor false. Either would be unfortunate. So, we will specify that our sentences have neither vice.

We can formulate our own revised version of the principle of bivalence, which states that:

> Principle of Bivalence: Each sentence of our language must be either true or false, not both, not neither.

This requirement may sound trivial, but in fact it constrains what we do from now on in interesting and even surprising ways. Even as we build more complex logical languages later, this principle will be fundamental.

Some readers may be thinking: what if I reject bivalence, or the principle of non-contradiction? There is a long line of philosophers who would like to argue with you, and propose that either move would be a mistake, and perhaps even incoherent. Set those arguments aside. If you have doubts about bivalence, or the principle of non-contradiction, stick with logic. That is because we could develop a logic in which there were more than two truth values. Logics have been created and studied in which we allow for three truth values, or continuous truth values, or stranger possibilities. The issue for us is that we must start somewhere, and the principle of bivalence is an intuitive way and—it would seem—the simplest way to start with respect to truth values. Learn basic logic first, and then you can explore these alternatives.

This points us to an important feature, and perhaps a mystery, of logic. In part, what a logical language shows us is the consequences of our assumptions. That might sound trivial, but, in fact, it is anything but. From very simple assumptions, we will discover new, and ultimately shocking, facts. So, if someone wants to study a logical language where we reject the principle of bivalence, they can do so. The difference between what they are doing, and what we will do in the following chapters, is that they will discover the consequences of rejecting the principle of bivalence, whereas we will discover the consequences of adhering to it. In either case, it would be wise to learn traditional logic first, before attempting to study or develop an alternative logic.

We should note at this point that we are not going to try to explain what "true" and "false" mean, other than saying that "false" means *not true.* When we add something to our language without explaining its meaning, we call it a "primitive". Philosophers have done much to try to understand what truth is, but it remains quite difficult to define truth in any way that is not controversial. Fortunately, taking *true* as a primitive will not get us into trouble, and it appears unlikely to make logic mysterious. We all have some grasp of what "true" means, and this grasp will be sufficient for our development of the propositional logic.

1.3 Atomic sentences

Our language will be concerned with declarative sentences, sentences that are either true or false, never both, and never neither. Here are some example sentences.

> 2+2=4.
>
> Malcolm Little is tall.
>
> If Lincoln wins the election, then Lincoln will be President.
>
> The Earth is not the center of the universe.

These are all declarative sentences. These all appear to satisfy our principle of bivalence. But they differ in important ways. The first two sentences do not have sentences as parts. For example, try to break up the first sentence. "2+2" is a function. "4" is a name. "=4" is a meaningless fragment, as is "2+". Only the whole expression, "2+2=4", is a sentence with a truth value. The second sentence is similar in this regard. "Malcolm Little" is a name. "is tall" is an adjective phrase (we will discover later that logicians call this a "predicate"). "Malcolm Little is" or "is tall" are fragments, they have no truth value.[2] Only "Malcolm Little is tall" is a complete sentence.

The first two example sentences above are of a kind we call "atomic sentences". The word "atom" comes from the ancient Greek word "atomos", meaning *cannot be cut.* When the ancient Greeks reasoned about matter, for example, some of them believed that if you took some substance, say a rock, and cut it into pieces, then cut the pieces into pieces, and so on, eventually you would get to something that could not be cut. This would be the smallest possible thing. (The fact that we now talk of having "split the atom" just goes to show that we changed the meaning of the word "atom". We came to use it as a name for a particular kind of thing, which then turned out to have parts, such as electrons, protons, and neutrons.) In logic, the idea of an atomic sentence is of a sentence that can have no parts that are sentences.

In reasoning about these atomic sentences, we could continue to use English. But for reasons that become clear as we proceed, there are many advantages to coming up with our own way of writing our sentences. It is traditional in logic to use upper case letters from **P** on (**P**, **Q**, **R**, **S**….) to stand for atomic sentences. Thus, instead of writing

Malcolm Little is tall.

We could write

P

If we want to know how to translate **P** to English, we can provide a translation key. Similarly, instead of writing

Malcolm Little is a great orator.

We could write

Q

And so on. Of course, written in this way, all we can see about such a sentence is that it is a sentence, and that perhaps **P** and **Q** are different sentences. But for now, these will be sufficient.

Note that not all sentences are atomic. The third sentence in our four examples above contains parts that are sentences. It contains the atomic sentence, "Lincoln wins the election" and also the atomic sentence, "Lincoln will be President". We could represent this whole sentence with a single letter. That is, we could let

If Lincoln wins the election, Lincoln will be president.

be represented in our logical language by

S

However, this would have the disadvantage that it would hide some of the sentences that are inside this sentence, and also it would hide their relationship. Our language would tell us more if we could capture the relation between the parts of this sentence, instead of hiding them. We will do this in chapter 2.

1.4 Syntax and semantics

An important and useful principle for understanding a language is the difference between syntax and semantics. "Syntax" refers to the "shape" of an expression in our language. It does not concern itself with what the elements of the language mean, but just specifies how they can be written out.

We can make a similar distinction (though not exactly the same) in a natural language. This expression in English has an uncertain meaning, but it has the right "shape" to be a sentence:

Colorless green ideas sleep furiously.

In other words, in English, this sentence is syntactically correct, although it may express some kind of meaning error.

An expression made with the parts of our language must have correct syntax in order for it to be a sentence. Sometimes, we also call an expression with the right syntactic form a "well-formed formula".

We contrast syntax with semantics. "Semantics" refers to the meaning of an expression of our language. Semantics depends upon the relation of that element of the language to something else. For example, the truth value of the sentence, "The Earth has one moon" depends not upon the English language, but upon something exterior to the language. Since the self-standing elements of our propositional logic are sentences, and the most important property of these is their truth value, the only semantic feature of sentences that will concern us in our propositional logic is their truth value.

Whenever we introduce a new element into the propositional logic, we will specify its syntax and its semantics. In the propositional logic, the syntax is generally trivial, but the semantics is less so. We have so far introduced atomic sentences. The syntax for an atomic sentence is trivial. If **P** is an atomic sentence, then it is syntactically correct to write down

P

By saying that this is syntactically correct, we are not saying that **P** is true. Rather, we are saying that **P** is a sentence.

If semantics in the propositional logic concerns only truth value, then we know that there are only two possible semantic values for **P**; it can be either true or false. We have a way of writing this that will later prove helpful. It is called a "truth table". For an atomic sentence, the truth table is trivial, but when we look at other kinds of sentences their truth tables will be more complex.

The idea of a truth table is to describe the conditions in which a sentence is true or false. We do this by identifying all the atomic sentences that compose that sentence. Then, on the left side, we stipulate all the possible truth values of these atomic sentences and write these out. On the right side, we then identify under what conditions the sentence (that is composed of the other atomic sentences) is true or false.

The idea is that the sentence on the right is dependent on the sentence(s) on the left. So the truth table is filled in like this:

Atomic sentence(s) that compose the dependent sentence on the right	Dependent sentence composed of the atomic sentences on the left
All possible combinations of truth values of the composing atomic sentences	Resulting truth values for each possible combination of truth values of the composing atomic sentences

We stipulate all the possible truth values on the bottom left because the propositional logic alone will not determine whether an atomic sentence is true or false; thus, we will simply have to consider both possibilities. Note that there are many ways that an atomic sentence can be true, and there are many ways that it can be false. For example, the sentence, "Tom is American" might be true if Tom was born in New York, in Texas, in Ohio, and so on. The sentence might be false because Tom was born to Italian parents in Italy, to French parents in France, and so on. So, we group all these cases together into two kinds of cases.

These are two rows of the truth table for an atomic sentence. Each row of the truth table represents a kind of way that the world could be. So here is the left side of a truth table with only a single atomic sentence, **P**. We will write "*T*" for *true* and "*F*" for *false.*

P	
T	
F	

There are only two relevant kinds of ways that the world can be, when we are considering the semantics of an atomic sentence. The world can be one of the many conditions such that **P** is true, or it can be one of the many conditions such that **P** is false.

To complete the truth table, we place the dependent sentence on the top right side, and describe its truth value in relation to the truth value of its parts. We want to identify the semantics of **P**, which has only one part, **P**. The truth table thus has the final form:

P	**P**
T	*T*
F	*F*

This truth table tells us the meaning of **P**, as far as our propositional logic can tell us about it. Thus, it gives us the complete semantics for **P**. (As we will see later, truth tables have three uses: to provide the semantics for a kind of sentence; to determine under what conditions a complex sentence is true or false; and to determine if an argument is good. Here we are describing only this first use.)

In this truth table, the first row combined together all the kinds of ways the world could be in which **P** is true. In the second column we see that for all of these kinds of ways the world could be in which **P** is true, unsurprisingly, **P** is true. The second row combines together all the kinds of ways the world could be in which **P** is false. In those, **P** is false. As we noted above, in the case of an atomic sentence, the truth table is trivial. Nonetheless, the basic concept is very useful, as we will begin to see in the next chapter.

One last tool will be helpful to us. Strictly speaking, what we have done above is give the syntax and semantics for a particular atomic sentence, **P**. We need a way to make general claims about all the sentences of our language, and then give the syntax and semantics for any atomic sentences. We do this using variables, and here we will use Greek letters for those variables, such as **Φ** and Ψ. Things said using these variables is called our "metalanguage", which means literally the *after language*, but which we take to mean, *our language about our language.* The particular propositional logic that we create is called our "object language". **P** and **Q** are sentences of our object language. **Φ** and Ψ are elements of our metalanguage. To specify now the syntax of atomic sentences (that is, of all atomic sentences) we can say: If **Φ** is an atomic sentence, then

Φ

is a sentence. This tells us that simply writing **Φ** down (whatever atomic sentence it may be), as we have just done, is to write down something that is syntactically correct.

To specify now the semantics of atomic sentences (that is, of all atomic sentences) we can say: If **Φ** is an atomic sentence, then the semantics of **Φ** is given by

Φ	**Φ**
T	*T*
F	*F*

Note an important and subtle point. The atomic sentences of our propositional logic will be what we call "contingent" sentences. A contingent sentence can be either true or false. We will see later that some complex sentences of our propositional logic must be true, and some complex sentences of our propositional logic must be false. But for the propositional logic, every atomic sentence is (as far as we can tell using the propositional logic alone) contingent. This observation matters because it greatly helps to clarify where logic begins, and where the methods of another discipline ends. For example, suppose we have an atomic sentence like:

Force is equal to mass times acceleration.

Igneous rocks formed under pressure.

Germany inflated its currency in 1923 in order to reduce its reparations debt.

Logic cannot tell us whether these are true or false. We will turn to physicists, and use their methods, to evaluate the first claim. We will turn to geologists, and use their methods, to evaluate the second claim. We will turn to historians, and use their methods, to evaluate the third claim. But the logician can tell the physicist, geologist, and historian what follows from their claims.

1.5 Problems

1. Vagueness arises when the conditions under which a sentence might be true are "fuzzy". That is, in some cases, we cannot identify if the sentence is true or false. If we

say, “Tom is tall”, this sentence is certainly true if Tom is the tallest person in the world, but it is not clear whether it is true if Tom is 185 centimeters tall. Identify or create five declarative sentences in English that are vague.

2. Ambiguity usually arises when a word or phrase has several distinct possible interpretations. In our example above, the word “pen” could mean either a writing implement or a structure to hold a child. A sentence that includes “pen” could be ambiguous, in which case it might be true for one interpretation and false for another. Identify or create five declarative sentences in English that are ambiguous. (This will probably require you to identify a homonym, a word that has more than one meaning but sounds or is written the same. If you are stumped, consider slang: many slang terms are ambiguous because they redefine existing words. For example, in the 1980s, in some communities and contexts, to say something was “bad” meant that it was good; this obviously can create ambiguous sentences.)

3. Often we can make a vague sentence precise by defining a specific interpretation of the meaning of an adjective, term, or other element of the language. For example, we could make the sentence “Tom is tall” precise by specifying one person referred to by “Tom”, and also by defining “…is tall” as true of anyone 180 centimeters tall or taller. For each of the five vague sentences that you identified or created for problem 1, describe how the interpretation of certain elements of the sentence could make the sentence no longer vague.

4. Often we can make an ambiguous sentence precise by specifying which of the possible meanings we intend to use. We could make the sentence, “Tom is by the pen” unambiguous by specifying which Tom we mean, and also defining “pen” to mean an infant play pen. For each of the five ambiguous sentences that you identified or created for problem 2, identify and describe how the interpretation of certain elements of the sentence could make the sentence no longer ambiguous.

5. Come up with five examples of your own of English sentences that are not declarative sentences. (Examples can include commands, exclamations, and promises.)

---

[2] There is a complex issue here that we will discuss later. But, in brief: “is” is ambiguous; it has several meanings. “Malcolm Little is” is a sentence if it is meant to assert the existence of Malcolm Little. The “is” that appears in the sentence, “Malcolm Little is tall”, however, is what we call the “‘is’ of predication”. In that sentence, “is” is used to assert that a property is had by Malcolm Little (the property of being tall); and here “is tall” is what we are calling a “predicate”. So, the “is” of predication has no clear meaning when appearing without the rest of the predicate; it does not assert existence.