A Concise Introduction to Logic

CRAIG DELANCEY

OPEN SUNY TEXTBOOKS

A Concise Introduction to Logic by Craig DeLancey is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License, except where otherwise noted.

© 2017 Craig DeLancey

ISBN: 978-1-942341-42-0 ebook
978-1-942341-43-7 print

This publication was made possible by a SUNY Innovative Instruction Technology Grant (IITG). IITG is a competitive grants program open to SUNY faculty and support staff across all disciplines. IITG encourages development of innovations that meet the Power of SUNY's transformative vision.

Published by Open SUNY Textbooks

Milne Library
State University of New York at Geneseo
Geneseo, NY 14454

Contents

PART I: PROPOSITIONAL LOGIC

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

$$\mathbf{\Phi}$$

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.)

6. Here are some sentences from literary works and other famous texts. Describe as best you can what the role of the sentence is. For example, the sentence might be a declarative sentence, which aims to describe things; or a question, which aims to solicit information; or a command, which is used to make someone do something; and so on. It is not essential that you have a name for the kind of sentence, but rather can you describe what a speaker would typically intend for such a sentence to do?

    a. "Though I should die with thee, yet will I not deny thee." (From the King James Bible)
    b. "Get thee to a nunnery." (William Shakespeare, *Hamlet*.)
    c. "That on the first day of January, in the year of our Lord one thousand eight hundred and sixty-three, all persons held as slaves within any State or designated part of a State, the people whereof shall then be in rebellion against the United States, shall be then, thenceforward, and

forever free." (From "The Emancipation Proclamation".)

d. "Sing, goddess, of the anger of Achilles son of Peleus, that brought countless ills upon the Achaeans." (Homer, *The Illiad*.)
e. "Since the heavens grant that you recognize me, hold your tongue, and do not say a word about who I am to any one else in the house, for if you do, and if heaven grants me to take the lives of these suitors, I will not spare you, though you are my own nurse, when I am killing the other women." (Homer, *The Odyssey*.)
f. "Tyger, tyger, burning bright,
In the forests of the night,
What immortal hand or eye,
could frame thy fearful symmetry?" (William Blake, *The Tyger*.)
g. "As wicked dew as e'er my mother brush'd
With raven's feather from unwholesome fen
Drop on you both! A south-west blow on ye
And blister you all o'er!" (William Shakespeare, *The Tempest*.)
h. "For he to-day that sheds his blood with me
Shall be my brother." (William Shakespeare, *Henry V*.)
i. "Astonishing, Pip!" (Charles Dickens, *Great Expectations*.)
j. "Congress shall make no law respecting an establishment of religion, or prohibiting the free exercise thereof; or abridging the freedom of speech, or of the press; or the right of the people peaceably to assemble, and to petition the government for a redress of grievances." (The Constitution of the United States.)

---

[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.

2. “If…then….” and “It is not the case that….”

2.1 The Conditional

As we noted in chapter 1, there are sentences of a natural language, like English, that are not atomic sentences. Our examples included

If Lincoln wins the election, then Lincoln will be President.

The Earth is not the center of the universe.

We could treat these like atomic sentences, but then we would lose a great deal of important information. For example, the first sentence tells us something about the relationship between the atomic sentences “Lincoln wins the election” and “Lincoln will be President”. And the second sentence above will, one supposes, have an interesting relationship to the sentence, “The Earth is the center of the universe”. To make these relations explicit, we will have to understand what “if…then…” and “not” mean. Thus, it would be useful if our logical language was able to express these kinds of sentences in a way that made these elements explicit. Let us start with the first one.

The sentence, “If Lincoln wins the election, then Lincoln will be President” contains two atomic sentences, “Lincoln wins the election” and “Lincoln will be President”. We could thus represent this sentence by letting

Lincoln wins the election

be represented in our logical language by

P

And by letting

Lincoln will be president

be represented by

Q

Then, the whole expression could be represented by writing

If **P** then **Q**

It will be useful, however, to replace the English phrase "if…then…" by a single symbol in our language. The most commonly used such symbol is "$\rightarrow$". Thus, we would write

P$\rightarrow$**Q**

One last thing needs to be observed, however. We might want to combine this complex sentence with other sentences. In that case, we need a way to identify that this is a single sentence when it is combined with other sentences. There are several ways to do this, but the most familiar (although not the most elegant) is to use parentheses. Thus, we will write our expression

(P$\rightarrow$**Q)**

This kind of sentence is called a "conditional". It is also sometimes called a "material conditional". The first constituent sentence (the one before the arrow, which in this example is "**P**") is called the "antecedent". The second sentence (the one after the arrow, which in this example is "**Q**") is called the "consequent".

We know how to write the conditional, but what does it mean? As before, we will take the meaning to be given by the truth conditions—that is, a description of when the sentence is either true or false. We do this with a truth table. But now, our sentence has two parts that are atomic sentences, **P** and **Q**. Note that either atomic sentence could be true or false. That means, we have to consider four possible kinds of situations. We must consider when **P** is true and when it is false, but then we need to consider those two kinds of situations twice: once for when **Q** is true and once for when **Q** is false. Thus, the left hand side of our truth table will look like this:

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

There are four kinds of ways the world could be that we must consider.

Note that, since there are two possible truth values (true and false), whenever we consider another atomic sentence, there are twice as many ways the world could be that we should consider. Thus, for n atomic sentences, our truth table must have 2^n rows. In the case of a conditional formed out of two atomic sentences, like our example of **(P→Q)**, our truth table will have 2^2 rows, which is 4 rows. We see this is the case above.

Now, we must decide upon what the conditional means. To some degree this is up to us. What matters is that once we define the semantics of the conditional, we stick to our definition. But we want to capture as much of the meaning of the English "if…then…" as we can, while remaining absolutely precise in our language.

Let us consider each kind of way the world could be. For the first row of the truth table, we have that **P** is true and **Q** is true. Suppose the world is such that Lincoln wins the election, and also Lincoln will be President. Then, would I have spoken truly if I said, "If Lincoln wins the election, then Lincoln will be President"? Most people agree that I would have. Similarly, suppose that Lincoln wins the election, but Lincoln will not be President. Would the sentence "If Lincoln wins the election, then Lincoln will be President" still be true? Most agree that it would be false now. So the first rows of our truth table are uncontroversial.

P	Q	(P→Q)
T	*T*	*T*
T	*F*	*F*
F	*T*	
F	*F*	

Some students, however, find it hard to determine what truth values should go in the next two rows. Note now that our principle of bivalence requires us to fill in these rows. We cannot leave them blank. If we did, we would be saying that sometimes a conditional can have no truth value; that is, we would be saying that sometimes, some sentences have no truth value. But our principle of bivalence requires that—in all kinds of situations—every sentence is either true or false, never both, never neither. So, if we are going to respect the principle of bivalence, then we have to put either ***T*** or ***F*** in for each of the last two rows.

It is helpful at this point to change our example. Let us consider two different examples to illustrate how best to fill out the remainder of the truth table for the conditional.

First, suppose I say the following to you: "If you give me \$50, then I will buy you a ticket to the concert tonight." Let

You give me \$50

be represented in our logic by

R

and let

I will buy you a ticket to the concert tonight.

be represented by

S

Our sentence then is

(R→S)

And its truth table—as far as we understand right now—is:

R	S	(R→S)
T	T	*T*
T	*F*	*F*
F	*T*	
F	*F*	

That is, if you give me the money and I buy you the ticket, my claim that "If you give me $50, then I will buy you a ticket to the concert tonight" is true. And, if you give me the money and I don't buy you the ticket, I lied, and my claim is false. But now, suppose you do not give me $50, but I buy you a ticket for the concert as a gift. Was my claim false? No. I simply bought you the ticket as a gift, but, presumably would have bought it if you gave me the money, also. Similarly, if you don't give me money, and I do not buy you a ticket, that seems perfectly consistent with my claim.

So, the best way to fill out the truth table is as follows.

R	S	(R→S)
T	*T*	*T*
T	*F*	*F*
F	*T*	*T*
F	*F*	*T*

Second, consider another sentence, which has the advantage that it is very clear with respect to these last two rows. Assume that ***a*** is a particular natural number, only you and I don't know what number it is (the natural numbers are the whole positive numbers: 1, 2, 3, 4…). Consider now the following sentence.

If ***a*** is evenly divisible by 4, then ***a*** is evenly divisible by 2.

(By "evenly divisible," I mean divisible without remainder.) The first thing to ask yourself is: is this sentence true? I hope we can all agree that it is—even though we do not know what ***a*** is. Let

a is evenly divisible by 4

be represented in our logic by

U

and let

a is evenly divisible by 2

be represented by

V

Our sentence then is

(U→V)

And its truth table—as far as we understand right now—is:

U	V	(U→V)
T	*T*	*T*
T	*F*	*F*
F	*T*	
F	*F*	

Now consider a case in which ***a*** is 6. This is like the third row of the truth table. It is not the case that 6 is evenly divisible by 4, but it is the case that 6 is evenly divisible by 2. And consider the case in which ***a*** is 7. This is like the fourth row of the truth table; 7 would be evenly divisible by neither 4 nor 2. But we agreed that the conditional is true—regardless of the value of ***a***! So, the truth table must be:[3]

U	V	(U→V)
T	*T*	*T*
T	*F*	*F*
F	*T*	*T*
F	*F*	*T*

Following this pattern, we should also fill out our table about the election with:

P	Q	(P→Q)
T	*T*	*T*
T	*F*	*F*
F	*T*	*T*
F	*F*	*T*

If you are dissatisfied by this, it might be helpful to think of these last two rows as vacuous cases. A conditional tells us about what happens if the antecedent is true. But when the antecedent is false, we simply default to true.

We are now ready to offer, in a more formal way, the syntax and semantics for the conditional.

The syntax of the conditional is that, if **Φ** and **Ψ** are sentences, then

(Φ→Ψ)

is a sentence.

The semantics of the conditional are given by a truth table. For any sentences **Φ** and **Ψ**:

Φ	**Ψ**	**(Φ→Ψ)**
T	*T*	*T*
T	*F*	*F*
F	*T*	*T*
F	*F*	*T*

Remember that this truth table is now a definition. It defines the meaning of "→". We are agreeing to use the symbol "→" to mean this from here on out.

The elements of the propositional logic, like "→", that we add to our language in order to form more complex sentences, are called "truth functional connectives". I hope it is clear why: the meaning of this symbol is given in a truth function. (If you are unfamiliar or uncertain about the idea of a function, think of a function as like a machine that takes in one or more inputs, and always then gives exactly one output. For the conditional, the inputs are two truth values; and the output is one truth value. For example, put *T F* into the truth function called "→", and you get out *F*.)

2.2 Alternative phrasings in English for the conditional. Only if.

English includes many alternative phrasings that appear to be equivalent to the conditional. Furthermore, in English and other natural languages, the order of the conditional will sometimes be reversed. We can capture the general sense of these cases by recognizing that each of the following phrasings would be translated as **(P→Q)**. (In these examples, we mix English and our propositional logic, in order to illustrate the variations succinctly.)

If **P**, then **Q**.

Q, if **P**.

On the condition that **P**, **Q**.

Q, on the condition that **P**.

Given that **P**, **Q**.

Q, given that **P**.

Provided that **P**, **Q**.

Q, provided that **P**.

When **P**, then **Q**.

Q, when **P**.

P implies **Q**.

Q is implied by **P**.

P is sufficient for **Q**.

Q is necessary for **P**.

An oddity of English is that the word "only" changes the meaning of "if". You can see this if you consider the following two sentences.

Fifi is a cat, if Fifi is a mammal.

Fifi is a cat only if Fifi is a mammal.

Suppose we know Fifi is an organism, but, we don't know what kind of organism Fifi is. Fifi could be a dog, a cat, a gray whale, a ladybug, a sponge. It seems clear that the first sentence is not necessarily true. If Fifi is a gray whale, for example, then it is true that Fifi is a mammal, but false that Fifi is a cat; and so, the first sentence would be false. But the second sentence looks like it must be true (given what you and I know about cats and mammals).

We should thus be careful to recognize that "only if" does not mean the same thing as "if". (If it did, these two sentences would have the same truth value in all situations.) In fact, it seems that "only if" can best be expressed by a conditional where the "only if" appears before the consequent (remember, the consequent is the second part of the conditional—the part that the arrows points at). Thus, sentences of this form:

P only if Q.

Only if Q, P.

are best expressed by the formula

(P→Q)

2.3 Test your understanding of the conditional

People sometimes find conditionals confusing. In part, this seems to be because some people confuse them with another kind of truth-functional connective, which we will learn about later, called the "biconditional". Also, sometimes "if…then…" is used in English in a different way (see section 17.7 if you are curious about alternative possible meanings). But from now on, we will understand the conditional as described above. To test whether you have properly grasped the conditional, consider the following puzzle.[4]

We have a set of four cards in figure 2.1. Each card has the following property: it has a shape on one side, and a letter on the other side. We shuffle and mix the cards, flipping some over while we shuffle. Then, we lay out the four cards:

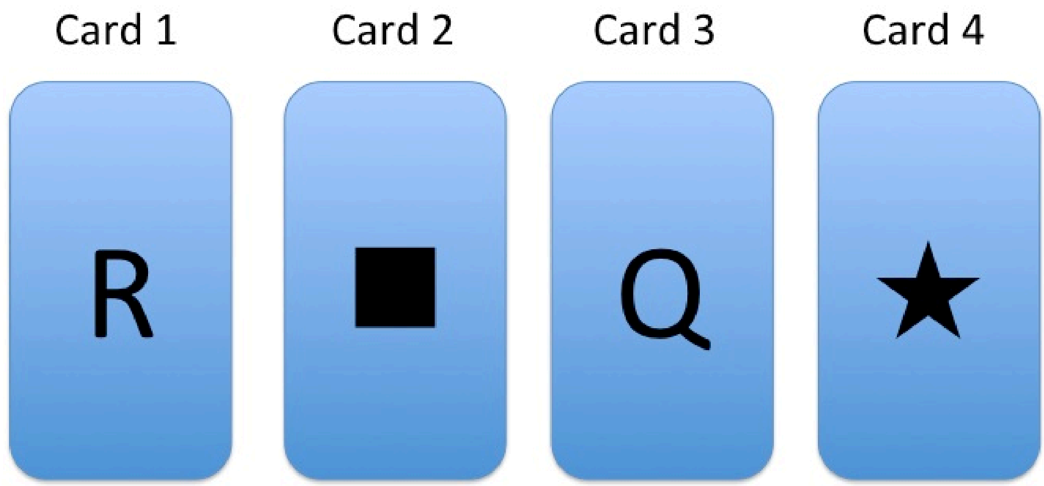

Figure 2.1

Given our constraint that each card has a letter on one side and a shape on the other, we know that card 1 has a shape on the unseen side; card 2 has a letter on the unseen side; and so on.

Consider now the following claim:

> For each of these four cards, if the card has a Q on the letter side of the card, then it has a square on the shape side of the card.

Here is our puzzle: what is the minimum number of cards that we must turn over to test whether this claim is true of all four cards; and which cards are they that we must turn over? Of course we could turn them all over, but the puzzle asks you to identify all and only the cards that will test the claim.

Stop reading now, and see if you can decide on the answer. Be warned, people generally perform poorly on this puzzle. Think about it for a while. The answer is given below in problem 1.

2.4 Alternative symbolizations for the conditional

Some logic books, and some logicians, use alternative symbolizations for the various truth-functional connectives. The meanings (that is, the truth tables) are always the same, but the symbol used may be different. For this reason, we will take the time in this text to briefly recognize alternative symbolizations.

The conditional is sometimes represented with the following symbol: "⊃". Thus, in such a case, $(P \rightarrow Q)$ would be written

$(P \supset Q)$

2.5 Negation

In chapter 1, we considered as an example the sentence,

The Earth is not the center of the universe.

At first glance, such a sentence might appear to be fundamentally unlike a conditional. It does not contain two sentences, but only one. There is a "not" in the sentence, but it is not connecting two sentences. However, we can still think of this sentence as being constructed with a truth functional connective, if we are willing to accept that this sentence is equivalent to the following sentence.

It is not the case that the Earth is the center of the universe.

If this sentence is equivalent to the one above, then we can treat "It is not the case" as a truth functional connective. It is traditional to replace this cumbersome English phrase with a single symbol, "¬". Then, mixing our propositional logic with English, we would have

¬The Earth is the center of the universe.

And if we let W be a sentence in our language that has the meaning *The Earth is the center of the universe*, we would write

$\neg W$

This connective is called "negation". Its syntax is: if Φ is a sentence, then

$$\neg\Phi$$

is a sentence. We call such a sentence a "negation sentence".

The semantics of a negation sentence is also obvious, and is given by the following truth table.

Φ	$\neg\Phi$
T	*F*
F	*T*

To deny a true sentence is to speak a falsehood. To deny a false sentence is to say something true.

Our syntax always is recursive. This means that syntactic rules can be applied repeatedly, to the product of the rule. In other words, our syntax tells us that if **P** is a sentence, then **¬P** is a sentence. But now note that the same rule applies again: if **¬P** is a sentence, then **¬¬P** is a sentence. And so on. Similarly, if **P** and **Q** are sentences, the syntax for the conditional tells us that **(P→Q)** is a sentence. But then so is **¬(P→Q)**, and so is **(¬(P→Q) → (P→Q))**. And so on. If we have just a single atomic sentence, our recursive syntax will allow us to form infinitely many different sentences with negation and the conditional.

2.6 Alternative symbolizations for negation

Some texts may use "~" for negation. Thus, **¬P** would be expressed with

~P

2.7 Problems

1. The answer to our card game was: you need only turn over cards 3 and 4. This might seem confusing to many people at first. But remember the

meaning of the conditional: it can only be false if the first part is true and the second part is false. The sentence we want to test is "For each of these four cards, if the card has a Q on the letter side of the card, then it has a square on the shape side of the card". Let **Q** stand for "the card has a Q on the letter side of the card." Let **S** stand for "the card has a square on the shape side of the card." Then we could make a truth table to express the meaning of the claim being tested:

Q	S	(Q→S)
T	*T*	*T*
T	*F*	*F*
F	*T*	*T*
F	*F*	*T*

Look back at the cards. The first card has an R on the letter side. So, sentence **Q** is false. But then we are in a situation like the last two rows of the truth table, and the conditional cannot be false. We do not need to check that card. The second card has a square on it. That means **S** is true for that card. But then we are in a situation represented by either the first or third row of the truth table. Again, the claim that **(Q→S)** cannot be false in either case with respect to that card, so there is no point in checking that card. The third card shows a Q. It corresponds to a situation that is like either the first or second row of the truth table. We cannot tell then whether **(Q→S)** is true or false of that card, without turning the card over. Similarly, the last card shows a situation where **S** is false, so we are in a kind of situation represented by either the second or last row of the truth table. We must turn the card over to determine if **(Q→S)** is true or false of that card.

Try this puzzle again. Consider the following claim about those same four cards: If there is a star on the shape side of the card, then there is an R on the letter side of the card. What is the minimum number of cards that you must turn over to check this claim? What cards are they?

2. Consider the following four cards in figure 2.2. Each card has a letter on one side, and a shape on the other side.

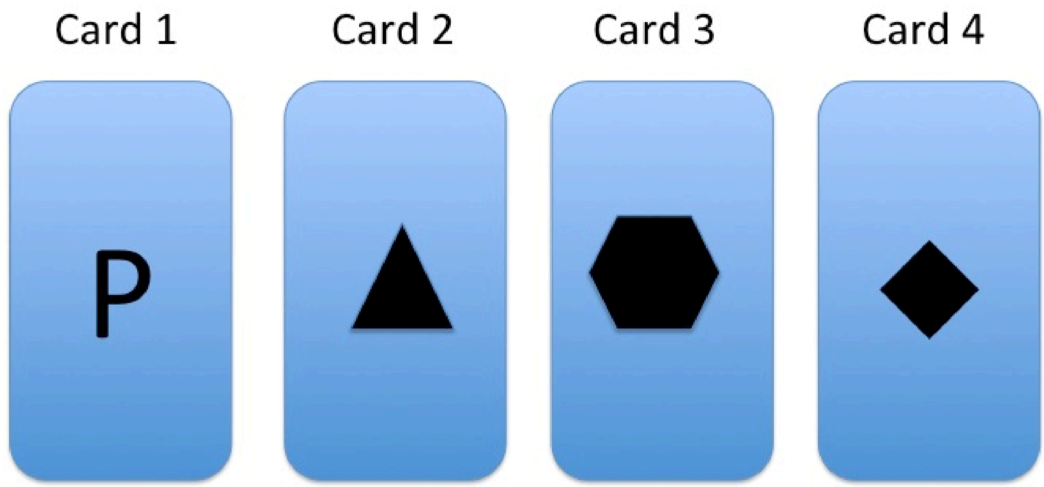

Figure 2.2

For each of the following claims, in order to determine if the claim is true of all four cards, describe (1) The minimum number of cards you must turn over to check the claim, and (2) what those cards are.

a. There is not a Q on the letter side of the card.
b. There is not an octagon on the shape side of the card.
c. If there is a triangle on the shape side of the card, then there is a P on the letter side of the card.
d. There is an R on the letter side of the card only if there is a diamond on the shape side of the card.
e. There is a hexagon on the shape side of the card, on the condition that there is a P on the letter side of the card.
f. There is a diamond on the shape side of the card only if there is a P on the letter side of the card.

3. Which of the following have correct syntax? Which have incorrect syntax?

a. $P \rightarrow Q$
b. $\neg(P \rightarrow Q)$
c. $(\neg P \rightarrow Q)$
d. $(P \neg \rightarrow Q)$
e. $(P \rightarrow \neg Q)$

f. $\neg\neg P$
g. $\neg P \neg$
h. $(\neg P \neg Q)$
i. $(\neg P \rightarrow \neg Q)$
j. $(\neg P \rightarrow \neg Q)\neg$

4. Use the following translation key to translate the following sentences into a propositional logic.

Translation Key	
Logic	English
P	Abe is able.
Q	Abe is honest.

a. If Abe is honest, Abe is able.
b. Abe is honest only if Abe is able.
c. Abe is able, if Abe is honest.
d. Only if Able is able, is Abe honest.
e. Abe is not able.
f. It's not the case that Abe isn't able.
g. Abe is not able only if Abe is not honest.
h. Abe is able, provided that Abe is not honest.
i. If Abe is not able then Abe is not honest.
j. It is not the case that, if Abe is able, then Abe is honest.

5. Make up your own translation key to translate the following sentences into a propositional logic. Then, use your key to translate the sentences into the propositional logic. Your translation key should contain only atomic sentences. These should be all and only the atomic sentences needed to translate the following sentences of English. Don't let it bother you that some of the sentences must be false.

a. Josie is a cat.
b. Josie is a mammal.
c. Josie is not a mammal.
d. If Josie is not a cat, then Josie is not a mammal.
e. Josie is a fish.
f. Provided that Josie is a mammal, Josie is not a fish.

g. Josie is a cat only if Josie is a mammal.
h. Josie is a fish only if Josie is not a mammal.
i. It's not the case that Josie is not a mammal.
j. Josie is not a cat, if Josie is a fish.

6. This problem will make use of the principle that our syntax is recursive. Translating these sentences is more challenging. Make up your own translation key to translate the following sentences into a propositional logic. Your translation key should contain only atomic sentences; these should be all and only the atomic sentences needed to translate the following sentences of English.

a. It is not the case that Tom won't pass the exam.
b. If Tom studies, Tom will pass the exam.
c. It is not the case that if Tom studies, then Tom will pass the exam.
d. If Tom does not study, then Tom will not pass the exam.
e. If Tom studies, Tom will pass the exam—provided that he wakes in time.
f. If Tom passes the exam, then if Steve studies, Steve will pass the exam.
g. It is not the case that if Tom passes the exam, then if Steve studies, Steve will pass the exam.
h. If Tom does not pass the exam, then if Steve studies, Steve will pass the exam.
i. If Tom does not pass the exam, then it is not the case that if Steve studies, Steve will pass the exam.
j. If Tom does not pass the exam, then if Steve does not study, Steve won't pass the exam.

7. Make up your own translation key in order to translate the following sentences into English. Write out the English equivalents in English sentences that seem (as much as is possible) natural.

a. $(\mathbf{R} \rightarrow \mathbf{S})$
b. $\neg\neg\mathbf{R}$
c. $(\mathbf{S} \rightarrow \mathbf{R})$
d. $\neg(\mathbf{S} \rightarrow \mathbf{R})$
e. $(\neg\mathbf{S} \rightarrow \neg\neg\mathbf{R})$
f. $\neg\neg(\mathbf{R} \rightarrow \mathbf{S})$
g. $(\neg\mathbf{R} \rightarrow \mathbf{S})$

h. **$(R \rightarrow \neg S)$**

i. **$(\neg R \rightarrow \neg S)$**

j. **$\neg(\neg R \rightarrow \neg S)$**

---

[3] One thing is a little funny about this second example with unknown number ***a***. We will not be able to find a number that is evenly divisible by 4 and not evenly divisible by 2, so the world will never be like the second row of this truth table describes. Two things need to be said about this. First, this oddity arises because of mathematical facts, not facts of our propositional logic—that is, we need to know what "divisible" means, what "4" and "2" mean, and so on, in order to understand the sentence. So, when we see that the second row is not possible, we are basing that on our knowledge of mathematics, not on our knowledge of propositional logic. Second, some conditionals can be false. In defining the conditional, we need to consider all possible conditionals; so, we must define the conditional for any case where the antecedent is true and the consequent is false, even if that cannot happen for this specific example.

[4] See Wason (1966).

3. Good Arguments

3.1 A historical example

An important example of excellent reasoning can be found in the case of the medical advances of the Nineteenth Century physician, Ignaz Semmelweis. Semmelweis was an obstetrician at the Vienna General Hospital. Built on the foundation of a poor house, and opened in 1784, the General Hospital is still operating today. Semmelweis, during his tenure as assistant to the head of one of two maternity clinics, noticed something very disturbing. The hospital had two clinics, separated only by a shared anteroom, known as the First and the Second Clinics. The mortality rate for mothers delivering babies in the First Clinic, however, was nearly three times as bad as the mortality for mothers in the Second Clinic (9.9 % average versus 3.4% average). The same was true for the babies born in the clinics: the mortality rate in the First Clinic was 6.1% versus 2.1% at the Second Clinic.[5] In nearly all these cases, the deaths were caused by what appeared to be the same illness, commonly called "childbed fever". Worse, these numbers actually understated the mortality rate of the First Clinic, because sometimes very ill patients were transferred to the general treatment portion of the hospital, and when they died, their death was counted as part of the mortality rate of the general hospital, not of the First Clinic.

Semmelweis set about trying to determine why the First Clinic had the higher mortality rate. He considered a number of hypotheses, many of which were suggested by or believed by other doctors.

One hypothesis was that cosmic-atmospheric-terrestrial influences caused childbed fever. The idea here was that some kind of feature of the atmosphere would cause the disease. But, Semmelweis observed, the First and Second Clinics were very close to each other, had similar ventilation, and shared a common anteroom. So, they had similar atmospheric conditions. He reasoned: If childbed fever is caused by cosmic-atmospheric-terrestrial influences, then the mortality rate would be similar in the First and Second Clinics. But the mortality rate was not similar in the First and Second Clinics. So, the childbed fever was not caused by cosmic-atmospheric-terrestrial influences.

Another hypothesis was that overcrowding caused the childbed fever. But, if overcrowding caused the childbed fever, then the more crowded of the two clinics should have the higher mortality rate. But, the Second Clinic was more crowded (in part because, aware of its lower mortality rate, mothers fought desperately to be put there instead of in the First Clinic). It did not have a higher mortality rate. So, the childbed fever was not caused by overcrowding.

Another hypothesis was that fear caused the childbed fever. In the Second Clinic, the priest delivering last rites could walk directly to a dying patient's room. For reasons of the layout of the rooms, the priest delivering last rites in the First Clinic walked by all the rooms, ringing a bell announcing his approach. This frightened patients; they could not tell if the priest was coming for them. Semmelweis arranged a different route for the priest and asked him to silence his bell. He reasoned: if the higher rate of childbed fever was caused by fear of death resulting from the priest's approach, then the rate of childbed fever should decline if people could not tell when the priest was coming to the Clinic. But it was not the case that the rate of childbed fever declined when people could not tell if the priest was coming to the First Clinic. So, the higher rate of childbed fever in the First Clinic was not caused by fear of death resulting from the priest's approach.

In the First Clinic, male doctors were trained; this was not true in the Second Clinic. These male doctors performed autopsies across the hall from the clinic, before delivering babies. Semmelweis knew of a doctor who cut himself while performing an autopsy, and who then died a terrible death not unlike that of the mothers who died of childbed fever. Semmelweis formed a hypothesis. The childbed fever was caused by something on the hands of the doctors, something that they picked up from corpses during autopsies, but that infected the women and infants. He reasoned that: if the fever was caused by cadaveric matter on the hands of the doctors, then the mortality rate would drop when doctors washed their hands with chlorinated water before delivering babies. He forced the doctors to do this. The result was that the mortality rate dropped to a rate below that even of the Second Clinic.

Semmelweis concluded that the best explanation of the higher mortality rate was this "cadaveric matter" on the hands of doctors. He was the first person to see that washing of hands with sterilizing cleaners would save thousands of lives. It is hard to overstate how important this contribution is to human well being. Semmelweis's fine reasoning deserves our endless respect and gratitude.

But how can we be sure his reasoning was good? Semmelweis was essentially considering a series of arguments. Let us turn to the question: how shall we evaluate arguments?

3.2 Arguments

Our logical language now allows us to say conditional and negation statements. That may not seem like much, but our language is now complex enough for us to develop the idea of using our logic not just to describe things, but also to reason about those things.

We will think of reasoning as providing an argument. Here, we use the word "argument" not in the sense of two or more people criticizing each other, but rather in the sense we mean when we say, "Pythagoras's argument". In such a case, someone is using language to try to convince us that something is true. Our goal is to make this notion very precise, and then identify what makes an argument good.

We need to begin by making the notion of an argument precise. Our logical language so far contains only sentences. An argument will, therefore, consist of sentences. In a natural language, we use the term "argument" in a strong way, which includes the suggestion that the argument should be good. However, we want to separate the notion of a good argument from the notion of an argument, so we can identify what makes an argument good, and what makes an argument bad. To do this, we will start with a minimal notion of what an argument is. Here is the simplest, most minimal notion:

> Argument: an ordered list of sentences; we call one of these sentences the "conclusion", and we call the other sentences "premises".

This is obviously very weak. (There is a famous Monty Python skit where one of the comedians ridicules the very idea that such a thing could be called an argument.) But for our purposes, this is a useful notion because it is very clearly defined, and we can now ask, what makes an argument good?

The everyday notion of an argument is that it is used to convince us to believe something. The thing that we are being encouraged to believe is the conclusion.

Following our definition of “argument”, the reasons that the person gives will be what we are calling “premises”. But *belief* is a psychological notion. We instead are interested only in truth. So, we can reformulate this intuitive notion of what an argument should do, and think of an argument as being used to show that something is true. The premises of the argument are meant to show us that the conclusion is true.

What then should be this relation between the premises and the conclusion? Intuitive notions include that the premises should support the conclusion, or corroborate the conclusion, or make the conclusion true. But “support” and “corroborate” sound rather weak, and “make” is not very clear. What we can use in their place is a stronger standard: let us say as a first approximation that if the premises are true, the conclusion is true.

But even this seems weak, on reflection. For, the conclusion could be true by accident, for reasons unrelated to our premises. Remember that we define the conditional as true if the antecedent and consequent are true. But this could happen by accident. For example, suppose I say, “If Tom wears blue then he will get an A on the exam”. Suppose also that Tom both wears blue and Tom gets an A on the exam. This makes the conditional true, but (we hope) the color of his clothes really had nothing to do with his performance on the exam. Just so, we want our definition of “good argument” to be such that it cannot be an accident that the premises and conclusion are both true.

A better and stronger standard would be that, necessarily, given true premises, the conclusion is true.

This points us to our definition of a good argument. It is traditional to call a good argument “valid.”

> Valid argument: an argument for which, necessarily, if the premises are true, then the conclusion is true.

This is the single most important principle in this book. Memorize it.

A bad argument is an argument that is not valid. Our name for this will be an “invalid argument”.

Sometimes, a dictionary or other book will define or describe a “valid argument” as an argument that follows the rules of logic. This is a hopeless way to define "valid", because it is circular in a pernicious way: we are going to create the rules of our logic in order to ensure that they construct valid arguments. We cannot make rules of logical reasoning until we know what we want those rules to do, and what we want them to do is to create valid arguments. So "valid" must be defined before we can make our reasoning system.

Experience shows that if a student is to err in understanding this definition of "valid argument", he or she will typically make the error of assuming that a valid argument has all true premises. This is not required. There are valid arguments with false premises and a false conclusion. Here's one:

> If Miami is the capital of Kansas, then Miami is in Canada. Miami is the capital of Kansas. Therefore, Miami is in Canada.

This argument has at least one false premise: Miami is not the capital of Kansas. And the conclusion is false: Miami is not in Canada. But the argument is valid: if the premises were both true, the conclusion would have to be true. (If that bothers you, hold on a while and we will convince you that this argument is valid because of its form alone. Also, keep in mind always that "if…then…" is interpreted as meaning the conditional.)

Similarly, there are invalid arguments with true premises, and with a true conclusion. Here's one:

> If Miami is the capital of Ontario, then Miami is in Canada. Miami is not the capital of Ontario. Therefore, Miami is not in Canada.

(If you find it confusing that this argument is invalid, look at it again after you finish reading this chapter.)

Validity is about the relationship between the sentences in the argument. It is not a claim that those sentences are true.

Another variation of this confusion seems to arise when we forgot to think carefully about the conditional. The definition of valid is not "All the premises are true, so the conclusion is true." If you don't see the difference, consider the following two sentences. "If your house is on fire, then you should call the fire department." In this sentence, there is no claim that your house is on fire. It is rather advice about what you should do if your house is on fire. In the same way, the definition of valid argument does not tell you that the premises are true. It tells you what follows if they are true. Contrast now, "Your house is on fire, so you should call the fire department". This sentence delivers very bad news. It is not a conditional at all. What it really means is, "Your house is on fire and you should call the fire department". Our definition of valid is not, "All the premises are true and the conclusion is true".

Finally, another common mistake is to confuse *true* and *valid*. In the sense that we are using these terms in this book, only sentences can be true or false, and only arguments can be valid and invalid. When discussing and using our logical language, it is nonsense to say, "a true argument", and it is nonsense to say, "a valid sentence".

Someone new to logic might wonder, why would we want a definition of "good argument" that does not guarantee that our conclusion is true? The answer is that logic is an enormously powerful tool for checking arguments, and we want to be able to identify what the good arguments are, independently of the particular premises that we use in the argument. For example, there are infinitely many particular arguments that have the same form as the valid argument given above. There are infinitely many particular arguments that have the same form as the invalid argument given above. Logic lets us embrace all the former arguments at once, and reject all those bad ones at once.

Furthermore, our propositional logic will not be able to tell us whether an atomic sentence is true. If our argument is about rocks, we must ask the geologist if the premises are true. If our argument is about history, we must ask the historian if the premises are true. If our argument is about music, we must ask the music theorist if the premises are true. But the logician can tell the geologist, the historian, and the musicologist whether her arguments are good or bad, independent of the particular premises.

We do have a common term for a good argument that has true premises. This is called "sound". It is a useful notion when we are applying our logic. Here is our definition:

> Sound argument: a valid argument with true premises.

A sound argument must have a true conclusion, given the definition of "valid".

3.3 Checking arguments semantically

Every element of our definition of "valid" is clear except for one. We know what "if…then…" means. We defined the semantics of the conditional in chapter 2. We have defined "argument", "premise", and "conclusion". We take *true* and *false* as primitives. But what does "necessarily" mean?

We define a valid argument as one where, necessarily, if the premises are true, then the conclusion is true. It would seem the best way to understand this is to say, there is no situation in which the premises are true but the conclusion is false. But then, what are these "situations"? Fortunately, we already have a tool that looks like it could help us: the truth table.

Remember that in the truth table, we put on the bottom left side all the possible combinations of truth values of some set of atomic sentences. Each row of the table then represents a kind of way the world could be. Using this as a way to understand "necessarily", we could rephrase our definition of valid to something like this, "In any kind of situation in which all the premises are true, the conclusion is true."

Let's try it out. We will need to use truth tables in a new way: to check an argument. That will require having not just one sentence, but several on the truth table. Consider an argument that looks like it should be valid.

> If Jupiter is more massive than Earth, then Jupiter has a stronger gravitational field than Earth. Jupiter is more massive than Earth. In conclusion, Jupiter has a stronger gravitational field than Earth.

This looks like it has the form of a valid argument, and it looks like an astrophysicist would tell us it is sound. Let's translate it to our logical language using the following translation key. (We've used up our letters, so I'm going to start over. We'll do that often: assume we are starting a new language each time we translate a new set of problems or each time we consider a new example.)

> P: Jupiter is more massive than Earth
>
> Q: Jupiter has a stronger gravitational field than Earth.

This way of writing out sentences of logic and sentences of English we can call a "translation key". We can use this format whenever we want to explain what our sentences mean in English.

Using this key, our argument would be formulated

$(P \rightarrow Q)$

P

______

Q

That short line is not part of our language, but rather is a handy tradition. When quickly writing down arguments, we write the premises, and then write the conclusion last, and draw a short line above the conclusion.

This is an argument: it is an ordered list of sentences, the first two of which are premises and the last of which is the conclusion.

To make a truth table, we identify all the atomic sentences that constitute these sentences. These are **P** and **Q**. There are four possible kinds of ways the world could be that matter to us then:

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

We'll write out the sentences, in the order of premises and then conclusion.

		premise	premise	conclusion
P	Q	(P→Q)	P	Q
T	*T*			
T	*F*			
F	*T*			
F	*F*			

Now we can fill in the columns for each sentence, identifying the truth value of the sentence for that kind of situation.

		premise	premise	conclusion
P	Q	(P→Q)	P	Q
T	*T*	*T*	*T*	*T*
T	*F*	*F*	*T*	*F*
F	*T*	*T*	*F*	*T*
F	*F*	*T*	*F*	*F*

We know how to fill in the column for the conditional because we can refer back to the truth table used to define the conditional, to determine what its truth value is when the first part and second part are true; and so on. **P** is true in those kinds of situations where **P** is true, and **P** is false in those kinds of situations where **P** is false. And the same is so for **Q**.

Now, consider all those kinds of ways the world could be such that all the premises are true. Only the first row of the truth table is one where all the premises are true. Note that the conclusion is true in that row. That means, in any kind of situation in which all the premises are true, the conclusion will be true. Or, equivalently: necessarily, if all the premises are true, then the conclusion is true.

		premise	premise	conclusion
P	Q	$(P\rightarrow Q)$	P	Q
T	*T*	*T*	*T*	*T*
T	*F*	*F*	*T*	*F*
F	*T*	*T*	*F*	*T*
F	*F*	*T*	*F*	*F*

Consider in contrast the second argument above, the invalid argument with all true premises and a true conclusion. We'll use the following translation key.

R: Miami is the capital of Ontario

S: Miami is in Canada

And our argument is thus

$(R\rightarrow S)$

$\neg R$

_____

$\neg S$

Here is the truth table.

		premise	premise	conclusion
R	S	$(R\rightarrow S)$	$\neg R$	$\neg S$
T	*T*	*T*	*F*	*F*
T	*F*	*F*	*F*	*T*
F	*T*	*T*	*T*	*F*
F	*F*	*T*	*T*	*T*

Note that there are two kinds of ways that the world could be in which all of our premises are true. These correspond to the third and fourth row of the truth table. But for the third row of the truth table, the premises are true but the conclusion is false. Yes, there is a kind of way the world could be in which all the premises are true and the conclusion is true; that is shown in the fourth row of the truth table. But we are not interested in identifying arguments that will have true conclusions if we are lucky. We are interested in valid arguments. This argument is invalid. There is a kind of way the world could be such that all the premises are true and the conclusion is false. We can highlight this.

		premise	premise	conclusion
R	S	(R→S)	¬R	¬S
T	*T*	*T*	*F*	*F*
T	*F*	*F*	*F*	*T*
F	*T*	*T*	*T*	*F*
F	*F*	*T*	*T*	*T*

Hopefully it becomes clear why we care about validity. Any argument of the form, **(P→Q)** and **P**, therefore **Q**, is valid. We do not have to know what **P** and **Q** mean to determine this. Similarly, any argument of the form, **(R→S)** and **¬R**, therefore **¬S**, is invalid. We do not have to know what **R** and **S** mean to determine this. So logic can be of equal use to the astronomer and the financier, the computer scientist or the sociologist.

3.4 Returning to our historical example

We described some (not all) of the hypotheses that Semmelweis tested when he tried to identify the cause of childbed fever, so that he could save thousands of women and infants. Let us symbolize these and consider his reasoning.

The first case we considered was one where he reasoned: If childbed fever is caused by cosmic-atmospheric-terrestrial influences, then the mortality rate would be similar in the First and Second Clinics. But the mortality rate was not similar in the First and Second Clinics. So, the childbed fever is not caused by cosmic-atmospheric-terrestrial influences.

Here is a key to symbolize the argument.

> **T:** Childbed fever is caused by cosmic-atmospheric-terrestrial influences.
>
> **U:** The mortality rate is similar in the First and Second Clinics.

This would mean the argument is:

$(T \rightarrow U)$

$\neg U$

———

$\neg T$

Is this argument valid? We can check using a truth table.

		premise	premise	conclusion
T	U	$(T \rightarrow U)$	$\neg U$	$\neg T$
T	*T*	*T*	*F*	*F*
T	*F*	*F*	*T*	*F*
F	*T*	*T*	*F*	*T*
F	*F*	*T*	*T*	*T*

The last row is the only row where all the premises are true. For this row, the conclusion is true. Thus, for all the kinds of ways the world could be in which the premises are true, the conclusion is also true. This is a valid argument. If we accept his premises, then we should accept that childbed fever was not caused by cosmic-atmospheric-terrestrial influences.

The second argument we considered was the concern that fear caused the higher mortality rates, particularly the fear of the priest coming to deliver last rites. Semmelweis reasoned that if the higher rate of childbed fever is caused by fear of death resulting from the priest's approach, then the rate of childbed fever should decline if people cannot discern when the priest is coming to the Clinic. Here is a key:

> **V:** the higher rate of childbed fever is caused by fear of death resulting from the priest's approach.
>
> **W:** the rate of childbed fever will decline if people cannot discern when the priest is coming to the Clinic.

But when Semmelweis had the priest silence his bell, and take a different route, so that patients could not discern that he was coming to the First Clinic, he found no difference in the mortality rate; the First Clinic remained far worse than the second clinic. He concluded that the higher rate of childbed fever was not caused by fear of death resulting from the priest's approach.

$(V \rightarrow W)$

$\neg W$

_____

$\neg V$

Is this argument valid? We can check using a truth table.

		premise	premise	conclusion
V	W	$(V \rightarrow W)$	$\neg W$	$\neg V$
T	*T*	*T*	*F*	*F*
T	*F*	*F*	*T*	*F*
F	*T*	*T*	*F*	*T*
F	*F*	*T*	*T*	*T*

Again, we see that Semmelweis's reasoning was good. He showed that it was not the case that the higher rate of childbed fever was caused by fear of death resulting from the Priest's approach.

What about Semmelweis's positive conclusion, that the higher mortality rate was caused by some contaminant from the corpses that doctors had autopsied just before they assisted in a delivery? To understand this step in his method, we need to reflect a moment on the scientific method and its relation to logic.

3.5 Other kinds of arguments 1: Scientific reasoning

Valid arguments, and the methods that we are developing, are sometimes called "deductive reasoning". This is the kind of reasoning in which necessarily our conclusions is true if our premises are true; these arguments can be shown to be good by way of our logical reasoning alone. There are other kinds of reasoning, and understanding this may help clarify the relation of logic to other endeavors. Two important, and closely related, alternatives to deductive reasoning are scientific reasoning and statistical generalizations. We'll discuss statistical generalizations in the next section.

Scientific method relies upon logic, but science is not reducible to logic: scientists do empirical research. That is, they examine and test phenomena in the world. This is a very important difference from pure logic. To understand how this difference results in a distinct method, let us review Semmelweis's important discovery.

The details and nature of scientific reasoning are somewhat controversial. I am going to provide here a basic—many philosophers would say, oversimplified—account of scientific reasoning. My goal is to indicate the relation between logic and the kind of reasoning Semmelweis may have used.

As we noted, Semmelweis learned about the death of a colleague, Professor Jakob Kolletschka. Kolletschka had been performing an autopsy, and he cut his finger. Shortly thereafter, Kolletschka died with symptoms like those of childbed fever. Semmelweis reasoned that something on the corpse caused the disease; he called this "cadaveric matter". In the First Clinic, where the mortality rate of women and babies was high, doctors were doing autopsies and then delivering babies immediately after. If he could get this cadaveric matter off the hands of the doctors, the rate of childbed fever should fall.

So, he reasoned thus: if the fever is caused by cadaveric matter on the hands of the doctors, then the mortality rate will drop when doctors wash their hands with chlorinated water before delivering babies. He forced the doctors to do this. The result was that the mortality rate dropped a very great deal, at times to below 1%.

Here is a key:

P: The fever is caused by cadaveric matter on the hands of the doctors.

Q: The mortality rate will drop when doctors wash their hands with chlorinated water before delivering babies.

And the argument appears to be something like this (as we will see, this isn't quite the right way to put it, but for now…):

(P→Q)

Q

_____

P

Is this argument valid? We can check using a truth table.

		premise	premise	conclusion
P	Q	(P→Q)	Q	P
T	*T*	*T*	*T*	*T*
T	*F*	*F*	*F*	*T*
F	*T*	*T*	*T*	*F*
F	*F*	*T*	*F*	*F*

From this, it looks like Semmelweis has used an invalid argument!

However, an important feature of scientific reasoning must be kept in mind. There is some controversy over the details of the scientific method, but the most basic view goes something like this. Scientists formulate hypotheses about the possible causes or features of a phenomenon. They make predictions based on these hypotheses, and then they perform experiments to test those predictions. The reasoning here uses the conditional: if the hypotheses is true, then the particular prediction will be true. If the experiment shows that the prediction is false, then the scientist rejects the hypothesis.[6] But if the prediction proved to be true, then the scientist has shown that the hypothesis may be true—at least, given the information we glean from the conditional and the consequent alone.

This is very important. Scientific conclusions are about the physical world, they are not about logic. This means that scientific claims are not necessarily true, in the sense of “necessarily” that we used in our definition of “valid”. Instead, science identifies claims that may be true, or (after some progress) are very likely to be true, or (after very much progress) are true.

Scientists keep testing their hypotheses, using different predictions and experiments. Very often, they have several competing hypotheses that have, so far, survived testing. To decide between these, they can use a range of criteria. In order of their importance, these include: choose the hypothesis with the most predictive power (the one that correctly predicts more kinds of phenomena); choose the hypothesis that will be most productive of other scientific theories; choose the hypothesis consistent with your other accepted hypotheses; choose the simplest hypothesis.

What Semmelweis showed was that it could be true that cadaveric matter caused the childbed fever. This hypothesis predicted more than any other hypothesis that the doctors had, and so for that reason alone this was the very best hypothesis. “But,” you might reason, “doesn’t that mean his conclusion was true? And don’t we know now, given all that we’ve learned, that his conclusion must be true?” No. He was far ahead of other doctors, and his deep insights were of great service to all of humankind. But the scientific method continued to refine Semmelweis’s ideas. For example, later doctors introduced the idea of microorganisms as the cause of childbed fever, and this refined and improved Semmelweis’s insights: it was not because the cadaveric matter came from corpses that it caused the disease; it was because the cadaveric matter contained particular micro-organisms that it caused the disease. So, further scientific progress showed his hypothesis could be revised and improved.

To review and summarize, with the scientific method:

1. We develop a hypothesis about the causes or nature of a phenomenon.
2. We predict what (hopefully unexpected) effects are a consequence of this hypothesis.
3. We check with experiments to see if these predictions come true:

- If the predictions prove false, we reject the hypothesis;[7]
- If the predictions prove true, we conclude that the hypothesis could be true. We continue to test the hypothesis by making other predictions (that is, we

return to step 2).

This means that a hypothesis that does not make testable predictions (that is, a hypothesis that cannot possibly be proven false) is not a scientific hypothesis. Such a hypothesis is called "unfalsifiable" and we reject it as unscientific.

This method can result in more than one hypothesis being shown to be possibly true. Then, we chose between competing hypotheses by using criteria like the following (here ordered by their relative importance; "theory" can be taken to mean a collection of one or more hypotheses):

1. Predictive power: the more that a hypothesis can successfully predict, the better it is.
2. Productivity: a hypothesis that suggests more new directions for research is to be preferred.
3. Coherence with Existing Theory: if two hypotheses predict the same amount and are equally productive, then the hypothesis that coheres with (does not contradict) other successful theories is preferable to one that does contradict them.
4. Simplicity: if two hypotheses are equally predictive, productive, and coherent with existing theories, then the simpler hypothesis is preferable.

Out of respect to Ignaz Semmelweis we should tell the rest of his story, although it means we must end on a sad note. Semmelweis's great accomplishment was not respected by his colleagues, who resented being told that their lack of hygiene was causing deaths. He lost his position at the First Clinic, and his successors stopped the program of washing hands in chlorinated water. The mortality rate leapt back to its catastrophically high levels. Countless women and children died. Semmelweis continued to promote his ideas, and this caused growing resentment. Eventually, several doctors in Vienna—not one of them a psychiatrist—secretly signed papers declaring Semmelweis insane. We do not know whether Semmelweis was mentally ill at this time. These doctors took him to an asylum on the pretense of having him visit in his capacity as a doctor; when he arrived, the guards seized Semmelweis. He struggled, and the guards at the asylum beat him severely, put him in a straightjacket, and left him alone in a locked room. Neglected in isolation, the wounds from his beating became infected and he died a week later.

It was years before Semmelweis's views became widely accepted and his accomplishment properly recognized. His life teaches many lessons, including unfortunately that even the most educated among us can be evil, petty, and willfully ignorant. Let us repay Semmelweis, as those in his own time did not, by remembering and praising his scientific acumen and humanity.

3.6 Other kinds of arguments 2: Statistical reasoning

Here we can say a few words about statistical generalizations—our goal being only to provide a contrast with deductive reasoning.

In one kind of statistical generalization, we have a population of some kind that we want to make general claims about. A population could be objects or events. So, a population can be a group of organisms, or a group of weather events. "Population" just means all the events or all the things we want to make a generalization about. Often however it is impossible to examine every object or event in the population, so what we do is gather a sample. A sample is some portion of the population. Our hope is that the sample is representative of the population: that whatever traits are shared by the members of the sample are also shared by the members of the population.

For a sample to representative, it must be random and large enough. "Random" in this context means that the sample was not chosen in any way that might distinguish members of the sample from the population, other than being members of the population. In other words, every member of the population was equally likely to be in the sample. "Large enough" is harder to define. Statisticians have formal models describing this, but suffice to say we should not generalize about a whole population using just a few members.

Here's an example. We wonder if all domestic dogs are descended from wolves. Suppose we have some genetic test to identify if an organism was a descendent of wolves. We cannot give the test to all domestic dogs—this would be impractical and costly and unnecessary. We pick a random sample of domestic dogs that is large enough, and we test them. For the sample to be random, we need to select it without allowing any bias to influence our selection; all that should matter is that these are domestic dogs, and each member of the population must have an equal chance of being in the sample. Consider the alternative: if we just tested one family of dogs—say, dogs that are large—we might end up selecting dogs that differed from others in a way that matters to our test. For example, maybe large dogs are descended from wolves, but small dogs are not. Other kinds of bias can creep in less obviously. We might just sample dogs in our local community, and it might just be that people in our community prefer large dogs, and again we would have a sample bias. So, we randomly select dogs, and give them the genetic test.

Suppose the results were positive. We reason that if all the members of the randomly selected and large enough sample (the tested dogs) have the trait, then it is very likely that all the members of the population (all dogs) have the trait. Thus: we could say that it appears very likely that all dogs have the trait. (This likelihood can be estimated, so that we can also sometimes say how likely it is that all members of the population have the trait.)

This kind of reasoning obviously differs from a deductive argument very substantially. It is a method of testing claims about the world, it requires observations, and its conclusion is likely instead of being certain.

But such reasoning is not unrelated to logic. Deductive reasoning is the foundation of these and all other forms of reasoning. If one must reason using statistics in this way, one relies upon deductive methods always at some point in one's arguments. There was a conditional at the penultimate step of our reasoning, for example (we said "if all the members of the randomly selected and large enough sample have the trait, then it is very likely that all the members of the population have the trait"). Furthermore, the foundations of these methods (the most fundamental descriptions of what these methods are) are given using logic and mathematics. Logic, therefore, can be seen as the study of the most fundamental form of reasoning, which will be used in turn by all other forms of reasoning, including scientific and statistical reasoning.

3.7 Problems

1. Make truth tables to show that the following arguments are valid. Circle or highlight the rows of the truth table that show the argument is valid (that is, all the rows where all the premises are true). Note that you will need eight

rows in the truth table for problems d-f, and sixteen rows in the truth table for problems g and h.

a. Premises: **(P→Q)**, ¬Q. Conclusion: ¬**P**.
b. Premises: ¬P. Conclusion: **(P→Q)**.
c. Premises: Q. Conclusion: **(P→Q)**.
d. Premises: **(P→Q)**, **(Q→R)**. Conclusion: **(P→R)**.
e. Premises: **(P→Q)**, **(Q→R)**, P. Conclusion: **R**.
f. Premises: **(P→Q)**, **(Q→R)**, ¬R. Conclusion: ¬**P**.
g. Premises: **(P→Q)**, **(Q→R)**, **(R→S)**, P. Conclusion: S.
h. Premises: **(P→Q)**, **(Q→R)**, **(R→S)**. Conclusion: **(P→S)**.

2. Make truth tables to show the following arguments are invalid. Circle or highlight the rows of the truth table that show the argument is invalid (that is, any row where all the premises are true but the conclusion is false).

   a. Premises: **(P→Q)**. Conclusion: **P**.
   b. Premises: **(P→Q)**. Conclusion: **Q**.
   c. Premises: **P**. Conclusion: **(P→Q)**.
   d. Premises: **(P→Q)**, Q. Conclusion: **P**.
   e. Premises: ¬Q. Conclusion: **(P→Q)**.
   f. Premises: **(P→Q)**. Conclusion: **(Q→P)**.
   g. Premises: **(P→Q)**, **(Q→R)**, ¬P. Conclusion: ¬**R**.
   h. Premises: **(P→Q)**, **(Q→R)**, R. Conclusion: **P**.
   i. Premises: **(P→Q)**, **(Q→R)**. Conclusion: **(R→P)**.
   j. Premises: **(P→Q)**, **(Q→R)**, **(R→S)**. Conclusion: **(S→P)**.

3. In normal colloquial English, write your own valid argument with at least two premises. Your argument should just be a paragraph (not an ordered list of sentences or anything else that looks like logic). Translate it into propositional logic and use a truth table to show it is valid.

4. In normal colloquial English, write your own invalid argument with at least two premises. Translate it into propositional logic and use a truth table to show it is invalid.

5. For each of the following, state whether the argument described could be: valid, invalid, sound, unsound.

a. An argument with false premises and a false conclusion.
b. An argument with true premises and a false conclusion.
c. An argument with false premises and a true conclusion.
d. An argument with true premises and a true conclusion.

---

[5] All the data cited here comes from Carter (1983) and additional biographical information comes from Carter and Carter (2008). These books are highly recommended to anyone interested in the history of science or medicine.

[6] It would be more accurate to say, if the prediction proves false, the scientist must reject either the hypothesis or some other premise of her reasoning. For example, her argument may include the implicit premise that her scientific instruments were operating correctly. She might instead reject this premise that her instruments are working correctly, change one of her instruments, and try again to test the hypothesis. See Duhem (1991). Or, to return to the case of Semmelweis, he might wonder whether he sufficiently established that there were no differences in the atmosphere between the two clinics; or he might wonder whether he sufficiently muffled the Priest's approach; or whether he recorded his results accurately; and so on. As noted, my account of scientific reasoning here is simplified.

[7] Or, as noted in note 6, we reject some other premise of the argument.

4. Proofs

4.1 A problem with semantic demonstrations of validity

Given that we can test an argument for validity, it might seem that we have a fully developed system to study arguments. However, there is a significant practical difficulty with our semantic method of checking arguments using truth tables (you may have already noted what this practical difficulty is, when you did problems 1e and 2e of chapter 3). Consider the following argument:

Alison will go to the party.

If Alison will go to the party, then Beatrice will.

If Beatrice will go to the party, then Cathy will.

If Cathy will go to the party, then Diane will.

If Diane will go to the party, then Elizabeth will.

If Elizabeth will go to the party, then Fran will.

If Fran will go to the party, then Giada will.

If Giada will go to the party, then Hilary will.

If Hillary will go to the party, then Io will.

If Io will go to the party, then Julie will.

_____

Julie will go to the party.

Most of us will agree that this argument is valid. It has a rather simple form, in which one sentence is related to the previous sentence, so that we can see the conclusion follows from the premises. Without bothering to make a translation key, we can see the argument has the following form.

P

$(P \rightarrow Q)$

$(Q \rightarrow R)$

$(R \rightarrow S)$

$(S \rightarrow T)$

$(T \rightarrow U)$

$(U \rightarrow V)$

$(V \rightarrow W)$

$(W \rightarrow X)$

$(X \rightarrow Y)$

_____

Y

However, if we are going to check this argument, then the truth table will require 1024 rows! This follows directly from our observation that for arguments or sentences composed of n atomic sentences, the truth table will require 2^n rows. This argument contains 10 atomic sentences. A truth table checking its validity must have 2^{10} rows, and 2^{10}=1024. Furthermore, it would be trivial to extend the argument for another, say, ten steps, but then the truth table that we make would require more than a million rows!

For this reason, and for several others (which become evident later, when we consider more advanced logic), it is very valuable to develop a syntactic proof method. That is, a way to check proofs not using a truth table, but rather using rules of syntax.

Here is the idea that we will pursue. A valid argument is an argument such that, necessarily, if the premises are true, then the conclusion is true. We will start just with our premises. We will set aside the conclusion, only to remember it as a goal. Then, we will aim to find a reliable way to introduce another sentence into the argument, with the special property that, if the premises are true, then this single additional sentence to the argument must also be true. If we could find a method to do that, and if after repeated applications of this method we were able to write down our conclusion, then we would know that, necessarily, if our premises are true then the conclusion is true.

The idea is more clear when we demonstrate it. The method for introducing new sentences will be called "inference rules". We introduce our first inference rules for the conditional. Remember the truth table for the conditional:

Φ	Ψ	$(\Phi \rightarrow \Psi)$
T	*T*	*T*
T	*F*	*F*
F	*T*	*T*
F	*F*	*T*

Look at this for a moment. If we have a conditional like $(P \rightarrow Q)$ (looking at the truth table above, remember that this would meant that we let Φ be P and Ψ be Q), do we know whether any other sentence is true? From $(P \rightarrow Q)$ alone we do not. Even if $(P \rightarrow Q)$ is true, P could be false or Q could be false. But what if we have some additional information? Suppose we have as premises both $(P \rightarrow Q)$ and P. Then, we would know that if those premises were true, Q must be true. We have already checked this with a truth table.

		premise	premise	
P	Q	$(P \rightarrow Q)$	P	Q
T	*T*	*T*	*T*	*T*
T	*F*	*F*	*T*	*F*
F	*T*	*T*	*F*	*T*
F	*F*	*T*	*F*	*F*

The first row of the truth table is the only row where all of the premises are true; and for it, we find that Q is true. This, of course, generalizes to any conditional. That is, we have that:

		premise	premise	
Φ	**Ψ**	**(Φ→Ψ)**	**Φ**	**Ψ**
T	*T*	*T*	*T*	*T*
T	*F*	*F*	*T*	*F*
F	*T*	*T*	*F*	*T*
F	*F*	*T*	*F*	*F*

We now capture this insight not using a truth table, but by introducing a rule. The rule we will write out like this:

$$
\begin{array}{l}
(\Phi \rightarrow \Psi) \\
\Phi \\
\hline
\Psi
\end{array}
$$

This is a syntactic rule. It is saying that whenever we have written down a formula in our language that has the shape of the first row (that is, whenever we have a conditional), and whenever we also have written down a formula that has the shape in the second row (that is, whenever we also have written down the antecedent of the conditional), then go ahead, whenever you like, and write down a formula like that in the third row (the consequent of the conditional). The rule talks about the shape of the formulas, not their meaning. But of course we justified the rule by looking at the meanings.

We describe this by saying that the third line is "derived" from the earlier two lines using the inference rule.

This inference rule is old. We are, therefore, stuck with its well-established, but not very enlightening, name: "modus ponens". Thus, we say, for the above example, that the third line is derived from the earlier two lines using modus ponens.

4.2 Direct proof

We need one more concept: that of a proof. Specifically, we'll start with the most fundamental kind of proof, which is called a "direct proof". The idea of a direct proof is: we write down as numbered lines the premises of our argument. Then, after this, we can write down any line that is justified by an application of an inference rule to earlier lines in the proof. When we write down our conclusion, we are done.

Let us make a proof of the simple argument above, which has premises **(P→Q)** and **P**, and conclusion **Q**. We start by writing down the premises and numbering them. There is a useful bit of notation that we can introduce at this point. It is known as a "Fitch bar", named after a logician Frederic Fitch, who developed this technique. We will write a vertical bar to the left, with a horizontal line indicating that the premises are above the line.

$$
\begin{array}{|l}
1.\ (\mathsf{P} \to \mathsf{Q}) \\
2.\ \mathsf{P} \\
\hline
\\
\end{array}
$$

It is also helpful to identify where these steps came from. We can do that with a little explanation written out to the right.

$$
\begin{array}{|ll}
1.\ (\mathsf{P} \to \mathsf{Q}) & \text{premise} \\
2.\ \mathsf{P} & \text{premise} \\
\hline
& \\
\end{array}
$$

Now, we are allowed to write down any line that follows from an earlier line using an inference rule.

$$
\begin{array}{|ll}
1.\ (\mathsf{P} \to \mathsf{Q}) & \text{premise} \\
2.\ \mathsf{P} & \text{premise} \\
\hline
3.\ \mathsf{Q} & \\
\end{array}
$$

And, finally, we want a reader to understand what rule we used, so we add that into our explanation, identifying the rule and the lines used.

1. $(P \rightarrow Q)$	premise
2. P	premise
3. Q	modus ponens, 1, 2

That is a complete direct proof.

Notice a few things. The numbering of each line, and the explanations to the right, are bookkeeping; they are not part of our argument, but rather are used to explain our argument. However, always do them because, it is hard to understand a proof without them. Also, note that our idea is that the inference rule can be applied to any earlier line, including lines themselves derived using inference rules. It is not just premises to which we can apply an inference rule. Finally, note that we have established that this argument must be valid. From the premises, and an inference rule that preserves validity, we have arrived at the conclusion. Necessarily, the conclusion is true, if the premises are true.

The long argument that we started the chapter with can now be given a direct proof.

1. P	premise
2. $(P \rightarrow Q)$	premise
3. $(Q \rightarrow R)$	premise
4. $(R \rightarrow S)$	premise
5. $(S \rightarrow T)$	premise
6. $(T \rightarrow U)$	premise
7. $(U \rightarrow V)$	premise
8. $(V \rightarrow W)$	premise
9. $(W \rightarrow X)$	premise
10. $(X \rightarrow Y)$	premise
11. Q	modus ponens, 2, 1
12. R	modus ponens, 3, 11
13. S	modus ponens, 4, 12
14. T	modus ponens, 5, 13
15. U	modus ponens, 6, 14
16. V	modus ponens, 7, 15
17. W	modus ponens, 8, 16
18. X	modus ponens, 9, 17
19. Y	modus ponens, 10, 18

From repeated applications of modus ponens, we arrived at the conclusion. If lines 1 through 10 are true, line 19 must be true. The argument is valid. And, we completed it with 19 steps, as opposed to writing out 1024 rows of a truth table.

We can see now one of the very important features of understanding the difference between syntax and semantics. Our goal is to make the syntax of our language perfectly mirror its semantics. By manipulating symbols, we manage to say something about the world. This is a strange fact, one that underlies one of the deeper possibilities of language, and also, ultimately, of computers.

4.3 Other inference rules

We can now introduce other inference rules. Looking at the truth table for the conditional again, what else do we observe? Many have noted that if the consequent of a conditional is false, and the conditional is true, then the antecedent of the conditional must be false. Written out as a semantic check on arguments, this will be:

		premise	premise	
Φ	Ψ	$(\Phi \rightarrow \Psi)$	$\neg\Psi$	$\neg\Phi$
T	*T*	*T*	*F*	*F*
T	*F*	*F*	*T*	*F*
F	*T*	*T*	*F*	*T*
F	*F*	*T*	*T*	*T*

(Remember how we have filled out the truth table. We referred to those truth tables used to define "→" and "¬", and then for each row of this table above, we filled out the values in each column based on that definition.)

What we observe from this truth table is that when both $(\Phi \rightarrow \Psi)$ and $\neg\Psi$ are true, then $\neg\Phi$ is true. Namely, this can be seen in the last row of the truth table.

This rule, like the last, is old, and has a well-established name: "modus tollens". We represent it schematically with

$$(\Phi \rightarrow \Psi)$$

$\neg\Psi$

_____

$\neg\Phi$

What about negation? If we know a sentence is false, then this fact alone does not tell us about any other sentence. But what if we consider a negated negation sentence? Such a sentence has the following truth table.

Φ	**¬¬Φ**
T	*T*
F	*F*

We can introduce a rule that takes advantage of this observation. In fact, it is traditional to introduce two rules, and lump them together under a common name. The rules' name is "double negation". Basically, the rule says we can add or take away two negations any time. Here are the two schemas for the two rules:

Φ

_____

$\neg\neg\Phi$

and

$\neg\neg\Phi$

_____

Φ

Finally, it is sometimes helpful to be able to repeat a line. Technically, this is an unnecessary rule, but if a proof gets long, we often find it easier to understand the proof if we write a line over again later when we find we need it again. So we introduce the rule "repeat".

Φ

$$\frac{}{\Phi}$$

4.4 An example

Here is an example that will make use of all three rules. Consider the following argument:

$(Q \to P)$

$(\neg Q \to R)$

$\neg R$

$\overline{}$

P

We want to check this argument to see if it is valid.

To do a direct proof, we number the premises so that we can refer to them when using inference rules.

1. $(Q \to P)$	premise
2. $(\neg Q \to R)$	premise
3. $\neg R$	premise

And, now, we apply our inference rules. Sometimes, it can be hard to see how to complete a proof. In the worst case, where you are uncertain of how to proceed, you can apply all the rules that you see are applicable and then, assess if you have gotten closer to the conclusion; and repeat this process. Here in any case is a direct proof of the sought conclusion.

1. $(Q \to P)$	premise
2. $(\neg Q \to R)$	premise
3. $\neg R$	premise
4. $\neg\neg Q$	modus tollens, 2, 3
5. Q	double negation, 4
6. P	modus ponens, 1, 5

Developing skill at completing proofs merely requires practice. You should strive to do as many problems as you can.

4.5 Problems

1. Complete a direct derivation (also called a "direct proof") for each of the following arguments, showing that it is valid. You will need the rules modus ponens, modus tollens, and double negation.

a. Premises: $(P \to Q)$, $\neg\neg P$. Show: $\neg\neg Q$.
b. Premises: Q, $(\neg P \to \neg Q)$. Show: P.
c. Premises: $\neg Q$, $(\neg Q \to S)$. Show: S.
d. Premises: $\neg S$, $(\neg Q \to S)$. Show: Q.
e. Premises: $(S \to \neg Q)$, $(P \to S)$, $\neg\neg P$. Show: $\neg Q$.
f. Premises: $(T \to P)$, $(Q \to S)$, $(S \to T)$, $\neg P$. Show: $\neg Q$.
g. Premises: R, P, $(P \to (R \to Q))$. Show: Q.
h. Premises: $((R \to S) \to Q)$, $\neg Q$, $(\neg(R \to S) \to V)$. Show: V.
i. Premises: $(P \to (Q \to R))$, $\neg(Q \to R)$. Show: $\neg P$.
j. Premises: $(\neg(Q \to R) \to P)$, $\neg P$, Q. Show: R.
k. Premises: P, $(P \to R)$, $(P \to (R \to Q))$. Show: Q.
l. Premises: $\neg R$, $(S \to R)$, P, $(P \to (T \to S))$. Show: $\neg T$.
m. Premises: P, $(P \to Q)$, $(P \to R)$, $(Q \to (R \to S))$. Show: S.
n. Premises: $(P \to (Q \to R))$, P, $((Q \to R) \to \neg S)$, $((T \to V) \to S)$. Show: $\neg(T \to V)$.

2. In normal colloquial English, write your own valid argument with at least two premises. Your argument should just be a paragraph (not an ordered list of sen-

tences or anything else that looks like logic). Translate it into propositional logic and use a direct proof to show it is valid.

3. In normal colloquial English, write your own valid argument with at least three premises. Your argument should just be a paragraph (not an ordered list of sentences or anything else that looks like logic). Translate it into propositional logic and use a direct proof to show it is valid.

4. Make your own key to translate into propositional logic the portions of the following argument that are in bold. Using a direct proof, prove that the resulting argument is valid.

> Inspector Tarski told his assistant, Mr. Carroll, "**If Wittgenstein had mud on his boots, then he was in the field**. Furthermore, **if Wittgenstein was in the field, then he is the prime suspect for the murder of Dodgson. Wittgenstein did have mud on his boots.** We conclude, **Wittgenstein is the prime suspect for the murder of Dodgson.**"

5. “And”

5.1 The conjunction

To make our logical language more easy and intuitive to use, we can now add to it elements that make it able to express the equivalents of other sentences from a natural language like English. Our translations will not be exact, but they will be close enough that: first, we will have a way to more quickly understand the language we are constructing; and, second, we will have a way to speak English more precisely when that is required of us.

Consider the following expressions. How would we translate them into our logical language?

> Tom will go to Berlin and Paris.
>
> The number a is evenly divisible by 2 and 3.
>
> Steve is from Texas but not from Dallas.

We could translate each of these using an atomic sentence. But then we would have lost—or rather we would have hidden—information that is clearly there in the English sentences. We can capture this information by introducing a new connective; one that corresponds to our “and”.

To see this, consider whether you will agree that these sentences above are equivalent to the following sentences.

> Tom will go to Berlin and Tom will go to Paris.
>
> The number a is evenly divisible by 2 and the number a is evenly divisible by 3.
>
> Steve is from Texas and it is not the case that Steve is from Dallas.

Once we grant that these sentences are equivalent to those above, we see that we can treat the “and” in each sentence as a truth functional connective.

Suppose we assume the following key.

P: Tom will go to Berlin.

Q: Tom will go to Paris.

R: a is evenly divisible by 2.

S: a is evenly divisible by 3.

T: Steve is from Texas

U: Steve is from Dallas.

A partial translation of these sentences would then be:

P and Q

R and S

T and ¬U

Our third sentence above might generate some controversy. How should we understand "but"? Consider that in terms of the truth value of the connected sentences, "but" is the same as "and". That is, if you say "P but Q" you are asserting that both P and Q are true. However, in English there is extra meaning; the English "but" seems to indicate that the additional sentence is unexpected or counter-intuitive. "P but Q" seems to say, "P is true, and you will find it surprising or unexpected that Q is true also." That extra meaning is lost in our logic. We will not be representing surprise or expectations. So, we can treat "but" as being the same as "and". This captures the truth value of the sentence formed using "but", which is all that we require of our logic.

Following our method up until now, we want a symbol to stand for "and". In recent years the most commonly used symbol has been "^".

The syntax for "^" is simple. If Φ and Ψ are sentences, then

(Φ^Ψ)

is a sentence. Our translations of our three example sentences should thus look like this:

(P^Q)

(R^S)

(T^¬U)

Each of these is called a "conjunction". The two parts of a conjunction are called "conjuncts".

The semantics of the conjunction are given by its truth table. Most people find the conjunction's semantics obvious. If I claim that both **Φ** and **Ψ** are true, normal usage requires that if **Φ** is false or **Ψ** is false, or both are false, then I spoke falsely also.

Consider an example. Suppose your employer says, "After one year of employment you will get a raise and two weeks vacation". A year passes. Suppose now that this employer gives you a raise but no vacation, or a vacation but no raise, or neither a raise nor a vacation. In each case, the employer has broken his promise. The sentence forming the promise turned out to be false.

Thus, the semantics for the conjunction are given with the following truth table. For any sentences **Φ** and **Ψ**:

Φ	**Ψ**	**(Φ^Ψ)**
T	*T*	*T*
T	*F*	*F*
F	*T*	*F*
F	*F*	*F*

5.2 Alternative phrasings, and a different "and"

We have noted that in English, "but" is an alternative to "and", and can be translated the same way in our propositional logic. There are other phrases that have a similar meaning: they are best translated by conjunctions, but they convey (in English) a sense of surprise or failure of expectations. For example, consider the following sentence.

Even though they lost the battle, they won the war.

Here “even though” seems to do the same work as “but”. The implication is that it is surprising—that one might expect that if they lost the battle then they lost the war. But, as we already noted, we will not capture expectations with our logic. So, we would take this sentence to be sufficiently equivalent to:

They lost the battle and they won the war.

With the exception of “but”, it seems in English there is no other single word that is an alternative to “and” that means the same thing. However, there are many ways that one can imply a conjunction. To see this, consider the following sentences.

Tom, who won the race, also won the championship.

The star Phosphorous, that we see in the morning, is the Evening Star.

The Evening Star, which is called “Hesperus”, is also the Morning Star.

While Steve is tall, Tom is not.

Dogs are vertebrate terrestrial mammals.

Depending on what elements we take as basic in our language, these sentences all include implied conjunctions. They are equivalent to the following sentences, for example:

Tom won the race and Tom won the championship.

Phosphorous is the star that we see in the morning and Phosphorous is the Evening Star.

The Evening Star is called “Hesperus” and the Evening Star is the Morning Star.

Steve is tall and it is not the case that Tom is tall.

Dogs are vertebrates and dogs are terrestrial and dogs are mammals.

Thus, we need to be sensitive to complex sentences that are conjunctions but that do not use “and” or “but” or phrases like “even though”.

Unfortunately, in English there are some uses of "and" that are not conjunctions. The same is true for equivalent terms in some other natural languages. Here is an example.

> Rochester is between Buffalo and Albany.

The "and" in this sentence is not a conjunction. To see this, note that this sentence is not equivalent to the following:

> Rochester is between Buffalo and Rochester is between Albany.

That sentence is not even semantically correct. What is happening in the original sentence?

The issue here is that "is between" is what we call a "predicate". We will learn about predicates in chapter 11, but what we can say here is that some predicates take several names in order to form a sentence. In English, if a predicate takes more than two names, then we typically use the "and" to combine names that are being described by that predicate. In contrast, the conjunction in our propositional logic only combines sentences. So, we must say that there are some uses of the English "and" that are not equivalent to our conjunction.

This could be confusing because sometimes in English we put "and" between names and there is an implied conjunction. Consider:

> Steve is older than Joe and Karen.

Superficially, this looks to have the same structure as "Rochester is between Buffalo and Albany". But this sentence really is equivalent to:

> Steve is older than Joe and Steve is older than Karen.

The difference, however, is that there must be three things in order for one to be between the other two. There need only be two things for one to be older than the other. So, in the sentence "Rochester is between Buffalo and Albany", we need all three names ("Rochester", "Buffalo", and "Albany) to make a single proper atomic sentence with "between". This tells us that the "and" is just being used to combine these names, and not to combine implied sentences (since there can be no implied sentence about what is "between", using just two or just one of these names).

That sounds complex. Do not despair, however. The use of “and” to identify names being used by predicates is less common than “and” being used for a conjunction. Also, after we discuss predicates in chapter 11, and after you have practiced translating different kinds of sentences, the distinction between these uses of “and” will become easy to identify in almost all cases. In the meantime, we shall pick examples that do not invite this confusion.

5.3 Inference rules for conjunctions

Looking at the truth table for the conjunction should tell us two things very clearly. First, if a conjunction is true, what else must be true? The obvious answer is that both of the parts, the conjuncts, must be true. We can introduce a rule to capture this insight. In fact, we can introduce two rules and call them by the same name, since the order of conjuncts does not affect their truth value. These rules are often called “simplification”.

$$\frac{(\Phi \wedge \Psi)}{\Phi}$$

And:

$$\frac{(\Phi \wedge \Psi)}{\Psi}$$

In other words, if $(\Phi \wedge \Psi)$ is true, then Φ must be true; and if $(\Phi \wedge \Psi)$ is true, then Ψ must be true.

We can also introduce a rule to show a conjunction, based on what we see from the truth table. That is, it is clear that there is only one kind of condition in which $(\Phi \wedge \Psi)$ is true, and that is when Φ is true and when Ψ is true. This suggests the following rule:

Φ

Ψ

———

(Φ^Ψ)

We might call this rule "conjunction", but to avoid confusion with the name of the sentences, we will call this rule "adjunction".

5.4 Reasoning with conjunctions

It would be helpful to consider some examples of reasoning with conjunctions. Let's begin with an argument in a natural language.

> Tom and Steve will go to London. If Steve goes to London, then he will ride the Eye. Tom will ride the Eye too, provided that he goes to London. So, both Steve and Tom will ride the Eye.

We need a translation key.

T: Tom will go to London.

S: Steve will go to London.

U: Tom will ride the Eye.

V: Steve will ride the Eye.

Thus our argument is:

(T^S)

(S→U)

(T→V)

———

(V^U)

Our direct proof will look like this.

1. $(\mathsf{T} \land \mathsf{S})$	premise
2. $(\mathsf{S} \rightarrow \mathsf{U})$	premise
3. $(\mathsf{T} \rightarrow \mathsf{V})$	premise
4. T	simplification, 1
5. V	modus ponens, 3, 4
6. S	simplification, 1
7. U	modus ponens, 2, 6
8. $(\mathsf{V} \land \mathsf{U})$	adjunction, 5, 7

Now an example using just our logical language. Consider the following argument.

(Q→¬S)

(P→(Q^R))

(T→¬R)

P

―――

(¬S^¬T)

Here is one possible proof.

1. $(\mathsf{Q} \rightarrow \neg\mathsf{S})$	premise
2. $(\mathsf{P} \rightarrow (\mathsf{Q} \land \mathsf{R}))$	premise
3. $(\mathsf{T} \rightarrow \neg\mathsf{R})$	premise
4. P	premise
5. $(\mathsf{Q} \land \mathsf{R})$	modus ponens, 2, 4
6. Q	simplification, 5
7. $\neg\mathsf{S}$	modus ponens, 1, 6
8. R	simplification, 5
9. $\neg\neg\mathsf{R}$	double negation, 8
10. $\neg\mathsf{T}$	modus tollens, 3, 9
11. $(\neg\mathsf{S} \land \neg\mathsf{T})$	adjunction, 7, 10

5.5 Alternative symbolizations for the conjunction

Alternative notations for the conjunction include the symbols "&" and the symbol "·". Thus, the expression **(P^Q)** would be written in these different styles, as:

(P&Q)

(P·Q)

5.6 Complex sentences

Now that we have three different connectives, this is a convenient time to consider complex sentences. The example that we just considered required us to symbolize complex sentences, which use several different kinds of connectives. We want to avoid confusion by being clear about the nature of these sentences. We also want to be able to understand when such sentences are true and when they are false. These two goals are closely related.

Consider the following sentences.

¬(P→Q)

(¬P→Q)

(¬P→¬Q)

We want to understand what kinds of sentences these are, and also when they are true and when they are false. (Sometimes people wrongly assume that there is some simple distribution law for negation and conditionals, so there is some additional value to reviewing these particular examples.) The first task is to determine what kinds of sentences these are. If the first symbol of your expression is a negation, then you know the sentence is a negation. The first sentence above is a negation. If the first symbol of your expression is a parenthesis, then for our logical language we know that we are dealing with a connective that combines two sentences.

The way to proceed is to match parentheses. Generally people are able to do this by eye, but if you are not, you can use the following rule. Moving left to right, the last "**(**" that you encounter always matches the first "**)**" that you encounter. These form a sentence that must have two parts combined with a connective. You can identify the two parts because each will be an atomic sentence, a negation sentence, or a more complex sentence bound with parentheses on each side of the connective.

In our propositional logic, each set of paired parentheses forms a sentence of its own. So, when we encounter a sentence that begins with a parenthesis, we find that if we match the other parentheses, we will ultimately end up with two sentences as constituents, one on each side of a single connective. The connective that combines these two parts is called the "main connective", and it tells us what kind of sentence this is. Thus, above we have examples of a negation, a conditional, and a conditional.

How should we understand the meaning of these sentences? Here we can use truth tables in a new, third way (along with defining a connective and checking arguments). Our method will be this.

First, write out the sentence on the right, leaving plenty of room. Identify what kind of sentence this is. If it is a negation sentence, you should add just to the left a column for the non-negated sentence. This is because the truth table defining negation tells us what a negated sentence means in relation to the non-negated sentence that forms the sentence. If the sentence is a conditional, make two columns to the left, one for the antecedent and one for the consequent. If the sentence is a conjunction, make two columns to the left, one for each conjunct. Here again, we do this because the semantic definitions of these connectives tell us what the truth value of the sentence is, as a function of the truth value of its two parts. Continue this process until the parts would be atomic sentences. Then, we stipulate all possible truth values for the atomic sentences. Once we have done this, we can fill out the truth table, working left to right.

Let's try it for $\neg(\mathbf{P}\rightarrow\mathbf{Q})$. We write it to the right.

¬(P→Q)

This is a negation sentence, so we write to the left the sentence being negated.

(P→Q)	¬(P→Q)

This sentence is a conditional. Its two parts are atomic sentences. We put these to the left of the dividing line, and we stipulate all possible combinations of truth values for these atomic sentences.

P	Q	(P→Q)	¬(P→Q)
T	*T*		
T	*F*		
F	*T*		
F	*F*		

Now, we can fill out each column, moving left to right. We have stipulated the values for **P** and **Q**, so we can identify the possible truth values of **(P→Q)**. The semantic definition for "→" tells us how to do that, given that we know for each row the truth value of its parts.

P	Q	(P→Q)	¬(P→Q)
T	*T*	*T*	
T	*F*	*F*	
F	*T*	*T*	
F	*F*	*T*	

This column now allows us to fill in the last column. The sentence in the last column is a negation of **(P→Q)**, so the definition of "¬" tell us that **¬(P→Q)** is true when **(P→Q)** is false, and **¬(P→Q)** is false when **(P→Q)** is true.

P	Q	(P→Q)	¬(P→Q)
T	*T*	*T*	*F*
T	*F*	*F*	*T*
F	*T*	*T*	*F*
F	*F*	*T*	F

This truth table tells us what **¬(P→Q)** means in our propositional logic. Namely, if we assert **¬(P→Q)** we are asserting that **P** is true and **Q** is false.

We can make similar truth tables for the other sentences.

P	Q	¬P	(¬P→Q)
T	*T*	*F*	*T*
T	*F*	*F*	*T*
F	*T*	*T*	*T*
F	*F*	*T*	*F*

How did we make this table? The sentence **(¬P→Q)** is a conditional with two parts, **¬P** and **Q**. Because **Q** is atomic, it will be on the left side. We make a row for **¬P**. The sentence **¬P** is a negation of **P**, which is atomic, so we put **P** also on the left. We fill in the columns, going left to right, using our definitions of the connectives.

And:

P	Q	¬P	¬Q	(¬P→¬Q)
T	*T*	*F*	*F*	*T*
T	*F*	*F*	*T*	*T*
F	*T*	*T*	*F*	*F*
F	*F*	*T*	*T*	*T*

Such a truth table is very helpful in determining when sentences are, and are not, equivalent. We have used the concept of equivalence repeatedly, but have not yet defined it. We can offer a semantic, and a syntactic, explanation of equivalence. The semantic notion is relevant here: we say two sentences **Φ** and **Ψ** are "equivalent" or "logically equivalent" when they must have the same truth value. (For the syntactic concept of equivalence, see section 9.2). These truth tables show that these three sentences are not equivalent, because it is not the case that they must have the same truth value. For example, if **P** and **Q** are both true, then **¬(P→Q)** is false but **(¬P→Q)** is true and **(¬P→¬Q)** is true. If **P** is false and **Q** is true, then **(¬P→Q)** is true but **(¬P→¬Q)** is false. Thus, each of these sentences is true in some situation where one of the others is false. No two of them are equivalent.

We should consider an example that uses conjunction, and which can help in some translations. How should we translate "Not both Steve and Tom will go to Berlin"? This sentence tells us that it is not the case that both Steve will go to Berlin and Tom will go to Berlin. The sentence does allow, however, that one of them will go to Berlin. Thus, let **U** mean *Steve will go to Berlin* and **V** mean *Tom will go to Berlin.* Then we should translate this sentence, **¬(U^V)**. We should not translate the sentence **(¬U^¬V)**. To see why, consider their truth tables.

U	V	(U^V)	¬(U^V)	¬U	¬V	(¬U^¬V)
T	*T*	*T*	*F*	*F*	*F*	*F*
T	*F*	*F*	*T*	*F*	*T*	*F*
F	*T*	*F*	*T*	*T*	*F*	*F*
F	*F*	*F*	*T*	*T*	*T*	*T*

We can see that **¬(U^V)** and **(¬U^¬V)** are not equivalent. Also, note the following. Both **¬(U^V)** and **(¬U^¬V)** are true if Steve does not go to Berlin and Tom does not go to Berlin. This is captured in the last row of this truth table, and this is consistent with the meaning of the English sentence. But, now note: it is true that not both Steve and Tom will go to Berlin, if Steve goes and Tom does not. This is captured in the second row of this truth table. It is true that not both Steve and Tom will go to Berlin, if Steve does not go but Tom does. This is captured in the third row of this truth table. In both kinds of cases (in both rows of the truth table), **¬(U^V)** is true but **(¬U^¬V)** is false. Thus, we can see that **¬(U^V)** is the correct translation of "Not both Steve and Tom will go to Berlin".

Let's consider a more complex sentence that uses all of our connectives so far: **((P^¬Q)→¬(P→Q))**. This sentence is a conditional. The antecedent is a conjunction. The consequent is a negation. Here is the truth table, completed.

P	Q	¬Q	(P→Q)	(P^¬Q)	¬(P→Q)	((P^¬Q)→¬(P→Q))
T	*T*	*F*	*T*	*F*	*F*	*T*
T	*F*	*T*	*F*	*T*	*T*	*T*
F	*T*	*F*	*T*	*F*	*F*	*T*
F	*F*	*T*	*T*	*F*	*F*	*T*

This sentence has an interesting property: it cannot be false. That is not surprising, once we think about what it says. In English, the sentence says: If **P** is true and **Q** is false, then it is not the case that **P** implies **Q**. That must be true: if it were the case that **P** implies **Q**, then if **P** is true then **Q** is true. But the antecedent says **P** is true and **Q** is false.

Sentences of the propositional logic that must be true are called "tautologies". We will discuss them at length in later chapters.

Finally, note that we can combine this method for finding the truth conditions for a complex sentence with our method for determining whether an argument is valid using a truth table. We will need to do this if any of our premises or the conclusion are complex. Here is an example. We'll start with an argument in English:

> If whales are mammals, then they have vestigial limbs. If whales are mammals, then they have a quadrupedal ancestor. Therefore, if whales are mammals then they have a quadrupedal ancestor and they have vestigial limbs.

We need a translation key.

P: Whales are mammals.

Q: Whales have have vestigial limbs.

R: Whales have a quadrupedal ancestor.

The argument will then be symbolized as:

(P→Q)

(P→R)

____

(P→(R^Q))

Here is a semantic check of the argument.

			premise	premise		conclusion
P	Q	R	(P→Q)	(P→R)	(R^Q)	(P→(R^Q))
T	*T*	*T*	*T*	*T*	*T*	*T*
T	*T*	*F*	*T*	*F*	*F*	*F*
T	*F*	*T*	*F*	*T*	*F*	*F*
T	*F*	*F*	*F*	*F*	*F*	*F*
F	*T*	*T*	*T*	*T*	*T*	*T*
F	*T*	*F*	*T*	*T*	*F*	*T*
F	*F*	*T*	*T*	*T*	*F*	*T*
F	*F*	*F*	*T*	*T*	*F*	*T*

We have highlighted the rows where the premises are all true. Note that for these, the conclusion is true. Thus, in any kind of situation in which all the premises are true, the conclusion is true. This is equivalent, we have noted, to our definition of valid: necessarily, if all the premises are true, the conclusion is true. So this is a valid argument. The third column of the analyzed sentences (the column for **(R^Q)**) is there so that we can identify when the conclusion is true. The conclusion is a conditional, and we needed to know, for each kind of situation, if its antecedent **P**, and if its consequent **(R^Q)**, are true. The third column tells us the situations in which the consequent is true. The stipulations on the left tell us in what kind of situation the antecedent P is true.

5.6 Problems

1. Translate the following sentences into our logical language. You will need to create your own key to do so.

    a. Ulysses, who is crafty, is from Ithaca.
    b. Ulysses, who isn't crafty, is from Ithaca.
    c. Ulysses, who is crafty, isn't from Ithaca.
    d. Ulysses isn't both crafty and from Ithaca.
    e. Ulysses will go home only if he's from Ithaca and not Troy.
    f. Ulysses is not both from Ithaca and Troy, though he is crafty.
    g. If Ulysses outsmarts both Circes and the Cyclops, then he can go home.
    h. If Ulysses outsmarts Circes but not the Cyclops, then he will be eaten.
    i. Though he won't outsmart Circe, Ulysses will outsmart the Cyclops, even given that he is from Ithaca.
    j. Ulysses won't outsmart both Circes and the Cyclops, but he won't be eaten and will go home even though he is from Ithaca.

2. Prove the following arguments are valid, using a direct derivation.

    a. Premise: $((P \to Q) \wedge \neg Q)$. Conclusion: $\neg P$.
    b. Premises: $(((P \to Q) \wedge (Q \to R)) \wedge P)$. Conclusion: R.
    c. Premises: $((P \to Q) \wedge (R \to S))$, $(\neg Q \wedge \neg S)$. Conclusion: $(\neg P \wedge \neg R)$.
    d. Premises: $((R \wedge S) \to T)$, $(Q \wedge \neg T)$. Conclusion: $\neg(R \wedge S)$.
    e. Premises: $(P \to (Q \to R))$, $(Q \wedge P)$. Conclusion: R.
    f. Premises: $(P \to (Q \to R))$, $(\neg R \wedge P)$. Conclusion: $\neg Q$.
    g. Premises: $((P \wedge Q) \to (R \to S))$, $(Q \wedge (P \wedge \neg S))$. Conclusion: $(\neg R \wedge Q)$.
    h. Premises: $((P \to Q) \wedge (R \to S))$, $(P \wedge R)$. Conclusion: $((P \wedge Q) \wedge (R \wedge S))$.

3. Make truth tables for the following complex sentences. Identify which are tautologies.

    a. $\neg(P \wedge Q)$

b. $\neg(\neg P \rightarrow \neg Q)$
c. $(P \wedge \neg P)$
d. $\neg(P \wedge \neg P)$
e. $(((P \rightarrow Q) \wedge \neg Q) \rightarrow \neg P)$
f. $(((P \rightarrow Q) \wedge \neg P) \rightarrow \neg Q)$
g. $(((P \rightarrow Q) \wedge P) \rightarrow Q)$
h. $(((P \rightarrow Q) \wedge Q) \rightarrow P)$

4. Make truth tables to show when the following sentences are true and when they are false. State which of these sentences are equivalent. Also, can you identify if any have the same truth table as some of our connectives?

    a. $\neg(P \wedge Q)$
    b. $(\neg P \wedge \neg Q)$
    c. $\neg(\neg P \wedge \neg Q)$
    d. $\neg(P \rightarrow Q)$
    e. $(P \wedge \neg Q)$
    f. $(\neg P \wedge Q)$
    g. $\neg(P \rightarrow \neg Q)$
    h. $\neg(\neg P \rightarrow \neg Q)$
    i. $(P \wedge (Q \wedge R))$
    j. $((P \wedge Q) \wedge R))$
    k. $(P \rightarrow (Q \rightarrow R))$
    l. $((P \rightarrow Q) \rightarrow R))$

5. Write a valid argument in normal colloquial English with at least two premises, one of which is a conjunction or includes a conjunction. Your argument should just be a paragraph (not an ordered list of sentences or anything else that looks like formal logic). Translate the argument into propositional logic. Prove it is valid.

6. Write a valid argument in normal colloquial English with at least three premises, one of which is a conjunction or includes a conjunction and one of which is a conditional or includes a conditional. Translate the argument into propositional logic. Prove it is valid.

7. Often in a natural language like English, there are many implicit conjunctions in descriptions and other phrases. Here are some passages from litera-

ture. Translate them into our propositional logic. You will want to make a separate key for each particular problem.

a. "But Achilles the son of Peleus again shouted at Agamemnon the son of Atreus, for he was still in a rage."
(Homer, *The Illiad*)
b. "Socrates is an evil-doer, and a curious person, who searches into things under the earth and in heaven, and he makes the worse appear the better cause…." (Plato, *The Apology*)
c. "Incensed with indignation, Satan stood
Unterrified…." (Milton, *Paradise Lost*)
d. "Teiresias, seer who comprehends all…
You know, though thy blind eyes see nothing,
What plague infects our city Thebes." (Sophocles, *Oedipus Rex*)
e. "Scrooge! a squeezing, wrenching, grasping, scraping, clutching, covetous, old sinner!" (Charles Dickens, "A Christmas Carrol")
f. "When I wrote the following pages, or rather the bulk of them, I lived alone, in the woods, a mile from any neighbor, in a house which I had built myself, on the shore of Walden Pond, in Concord, Massachusetts, and earned my living by the labor of my hands only." [Here one can substitute "Thoreau" for "I" in the translation, if helpful.]. (Henry David Thoreau, *Walden*)
g. "In appearance Shatov was in complete harmony with his convictions: he was short, awkward, had a shock of flaxen hair, broad shoulders, thick lips, very thick overhanging white eyebrows, a wrinkled forehead, and a hostile, obstinately downcast, as it were shamefaced, expression in his eyes." (Fyodor Dostoevsky, *The Possessed*)

8. Make your own key to translate the following argument into our propositional logic. Translate only the parts in bold. Prove the argument is valid.

> "I suspect Dr. Kronecker of the crime of stealing Cantor's book," Inspector Tarski said. His assistant, Mr. Carroll, waited patiently for his reasoning. "For," Tarski said, "The thief left cigarette ashes on the table. The thief also did not wear shoes, but slipped silently into the room. Thus, **If Dr. Kronecker smokes and is in his stocking feet, then he most likely stole Cantor's book.**" At this point, Tarski pointed at Kronecker's feet. "**Dr. Kronecker is in his stocking feet.**" Tarski reached forward and pulled

from Kronecker's pocket a gold cigarette case. "And **Kronecker smokes.**"

Mr. Carroll nodded sagely, "Your conclusion is obvious: **Dr. Kronecker most likely stole Cantor's book."**

6. Conditional Derivations

6.1 An argument from Hobbes

In his great work, *Leviathan*, the philosopher Thomas Hobbes (1588-1679) gives an important argument for government. Hobbes begins by claiming that without a common power, our condition is very poor indeed. He calls this state without government, “the state of nature”, and claims

> Hereby it is manifest that during the time men live without a common power to keep them all in awe, they are in that condition which is called war; and such a war as is of every man against every man…. In such condition there is no place for industry, because the fruit thereof is uncertain: and consequently no culture of the earth; no navigation, nor use of the commodities that may be imported by sea; no commodious building; no instruments of moving and removing such things as require much force; no knowledge of the face of the earth; no account of time; no arts; no letters; no society; and which is worst of all, continual fear, and danger of violent death; and the life of man, solitary, poor, nasty, brutish, and short.[8]

Hobbes developed what is sometimes called “contract theory”. This is a view of government in which one views the state as the product of a rational contract. Although we inherit our government, the idea is that in some sense we would find it rational to choose the government, were we ever in the position to do so. So, in the passage above, Hobbes claims that in this state of nature, we have absolute freedom, but this leads to universal struggle between all people. There can be no property, for example, if there is no power to enforce property rights. You are free to take other people’s things, but they are also free to take yours. Only violence can discourage such theft. But, a common power, like a king, can enforce rules, such as property rights. To have this common power, we must give up some freedoms. You are (or should be, if it were ever up to you) willing to give up those freedoms because of the benefits that you get from this. For example, you are willing to give up the freedom to just seize people’s goods, because you like even more that other people cannot seize your goods.

We can reconstruct Hobbes's defense of government, greatly simplified, as being something like this:

> If we want to be safe, then we should have a state that can protect us.
>
> If we should have a state that can protect us, then we should give up some freedoms.
>
> Therefore, if we want to be safe, then we should give up some freedoms.

Let us use the following translation key.

> **P:** We want to be safe.
>
> **Q:** We should have a state that can protect us.
>
> **R:** We should give up some freedoms.

The argument in our logical language would then be:

> $(P\rightarrow Q)$
>
> $(Q\rightarrow R)$
>
> _____
>
> $(P\rightarrow R)$

This is a valid argument. Let's take the time to show this with a truth table.

			premise	premise	conclusion
P	Q	R	$(P\rightarrow Q)$	$(Q\rightarrow R)$	$(P\rightarrow R)$
T	*T*	*T*	*T*	*T*	*T*
T	*T*	*F*	*T*	*F*	*F*
T	*F*	*T*	*F*	*T*	*T*
T	*F*	*F*	*F*	*T*	*F*
F	*T*	*T*	*T*	*T*	*T*
F	*T*	*F*	*T*	*F*	*T*
F	*F*	*T*	*T*	*T*	*T*
F	*F*	*F*	*T*	*T*	*T*

The rows in which all the premises are true are the first, fifth, seventh, and eighth rows. Note that in each such row, the conclusion is true. Thus, in any kind of situation where the premises are true, the conclusion is true. This is our semantics for a valid argument.

What syntactic method can we use to prove this argument is valid? Right now, we have none. Other than double negation, we cannot even apply any of our inference rules using these premises.

Some logic systems introduce a rule to capture this inference; this rule is typically called the "chain rule". But, there is a more general principle at stake here: we need a way to show conditionals. So we want to take another approach to showing this argument is valid.

6.2 Conditional derivation

As a handy rule of thumb, we can think of the inference rules as providing a way to either show a kind of sentence, or to make use of a kind of sentence. For example, adjunction allows us to show a conjunction. Simplification allows us to make use of a conjunction. But this pattern is not complete: we have rules to make use of a conditional (modus ponens and modus tollens), but no rule to show a conditional.

We will want to have some means to prove a conditional, because sometimes an argument will have a conditional as a conclusion. It is not clear what rule we should introduce, however. The conditional is true when the antecedent is false, or if both the antecedent and the consequent are true. That's a rather messy affair for making an inference rule.

However, think about what the conditional asserts: if the antecedent is true, then the consequent is true. We can make use of this idea not with an inference rule, but rather in the very structure of a proof. We treat the proof as embodying a conditional relationship.

Our idea is this: let us assume some sentence, Φ. If we can then prove another sentence Ψ, we will have proved that if Φ is true then Ψ is true. The proof structure will thus have a shape like this:

$$
\begin{array}{r|l}
 & \quad \phi \\
 & \quad \vdots \\
 & \quad \psi \\
\triangleright & (\phi \to \psi)
\end{array}
$$

The last line of the proof is justified by the shape of the proof: by assuming that **Φ** is true, and then using our inference rules to prove **Ψ**, we know that if **Φ** is true then **Ψ** is true. And this is just what the conditional asserts.

This method is sometimes referred to as an application of the deduction theorem. In chapter 17 we will prove the deduction theorem. Here, instead, we shall think of this as a proof method, traditionally called "conditional derivation".

A conditional derivation is like a direct derivation, but with two differences. First, along with the premises, you get a single special assumption, called "the assumption for conditional derivation". Second, you do not aim to show your conclusion, but rather the consequent of your conclusion. So, to show **(Φ→Ψ)** you will always assume **Φ** and try to show **Ψ**. Also, in our logical system, a conditional derivation will always be a subproof. A subproof is a proof within another proof. We always start with a direct proof, and then do the conditional proof within that direct proof.

Here is how we would apply the proof method to prove the validity of Hobbes's argument, as we reconstructed it above.

1. $(\mathsf{P} \to \mathsf{Q})$	premise
2. $(\mathsf{Q} \to \mathsf{R}))$	premise
3. P	assumption for conditional derivation
4 Q	modus ponens, 1, 3
5. R	modus ponens, 2, 4
6. $(\mathsf{P} \to \mathsf{R})$	conditional derivation, 3-5

Our Fitch bars make clear what is a sub-proof here; they let us see this as a direct derivation with a conditional derivation embedded in it. This is an important concept: we can have proofs within proofs.

An important principle is that once a subproof is done, we cannot use any of the lines in the subproof. We need this rule because conditional derivation allowed us to make a special assumption that we use only temporarily. Above, we assumed **P**. Our goal is only to show that if **P** is true, then **R** is true. But perhaps **P** isn't true. We do not want to later make use of **P** for some other purpose. So, we have the rule that when a subproof is complete, you cannot use the lines that occur in the subproof. In this case, that means that we cannot use lines 3, 4, or 5 for any other purpose than to show the conditional **(P→R)**. We cannot now cite those individual lines again. We can, however, use line 6, the conclusion of the subproof.

The Fitch bars—which we have used before now in our proofs only to separate the premises from the later steps—now have a very beneficial use. They allow us to set aside a conditional derivation as a subproof, and they help remind us that we cannot cite the lines in that subproof once the subproof is complete.

It might be helpful to give an example of why this is necessary. That is, it might be helpful to give an example of an argument made invalid because it makes use of lines in a finished subproof. Consider the following argument.

> If you are Pope, then you have a home in the Vatican.
>
> If you have a home in the Vatican, then you hear church bells often.
>
> _____
>
> If you are Pope, then you hear church bells often.

That is a valid argument, with the same form as the argument we adopted from Hobbes. However, if we broke our rule about conditional derivations, we could prove that you are Pope. Let's use this key:

> **S**: You are Pope.
>
> **T**: You have a home in the Vatican.
>
> **U**: You hear church bells often.

Now consider this "proof":

1. $(S \rightarrow T)$	premise
2. $(T \rightarrow U))$	premise
3. S	assumption for conditional derivation
4 T	modus ponens, 1, 3
5. U	modus ponens, 2, 4
6. $(S \rightarrow U)$	conditional derivation, 3-5
7. S	repeat, 3

And, thus, we have proven that you are Pope. But, of course, you are not the Pope. From true premises, we ended up with a false conclusion, so the argument is obviously invalid. What went wrong? The problem was that after we completed the conditional derivation that occurs in lines 3 through 5, and used that conditional derivation to assert line 6, we can no longer use those lines 3 through 5. But on line 7 we made use of line 3. Line 3 is not something we know to be true; our reasoning from lines 3 through line 5 was to ask, if S were true, what else would be true? When we are done with that conditional derivation, we can use only the conditional that we derived, and not the steps used in the conditional derivation.

6.3 Some additional examples

Here are a few kinds of arguments that help illustrate the power of the conditional derivation.

This argument makes use of conjunctions.

(P→Q)

(R→S)

_____

((P^R)→(Q^S))

We always begin by constructing a direct proof, using the Fitch bar to identify the premises of our argument, if any.

1. $(P \rightarrow Q)$	premise
2. $(R \rightarrow S)$	premise

Because the conclusion is a conditional, we assume the antecedent and show the consequent.

1. $(P \rightarrow Q)$	premise
2. $(R \rightarrow S)$	premise
3. $(P \wedge R)$	assumption for conditional derivation
4 P	simplification, 3
5. Q	modus ponens, 1, 4
6. R	simplification, 3
7. S	modus ponens, 2, 6
8. $(Q \wedge S)$	adjunction, 5, 7
9. $((P \wedge R) \rightarrow (Q \wedge S))$	conditional derivation, 3-8

Here's another example. Note that the following argument is valid.

$(P \rightarrow (S \rightarrow R))$

$(P \rightarrow (Q \rightarrow S))$

_____

$(P \rightarrow (Q \rightarrow R))$

The proof will require several embedded subproofs.

1. $(P \rightarrow (S \rightarrow R))$	premise
2. $(P \rightarrow (Q \rightarrow S))$	premise
3. P	assumption for conditional derivation
4. Q	assumption for conditional derivation
5. $(Q \rightarrow S)$	modus ponens, 2, 3
6. S	modus ponens, 5, 4
7. $(S \rightarrow R)$	modus ponens, 1, 3
8. R	modus ponens, 7, 6
9. $(Q \rightarrow R)$	conditional derivation, 4-8
10. $(P \rightarrow (Q \rightarrow R))$	conditional derivation, 3-9

6.4 Theorems

Conditional derivation allows us to see an important new concept. Consider the following sentence:

$$((P \rightarrow Q) \rightarrow (\neg Q \rightarrow \neg P))$$

This sentence is a tautology. To check this, we can make its truth table.

P	Q	¬Q	¬P	(P→Q)	(¬Q→¬P)	((P→Q) → (¬Q→¬P))
T	*T*	*F*	*F*	*T*	*T*	*T*
T	*F*	*T*	*F*	*F*	*F*	*T*
F	*T*	*F*	*T*	*T*	*T*	*T*
F	*F*	*T*	*T*	*T*	*T*	*T*

This sentence is true in every kind of situation, which is what we mean by a "tautology".

Now reflect on our definition of "valid": necessarily, if the premises are true, then the conclusion is true. What about an argument in which the conclusion is a tautology? By our definition of "valid", an argument with a conclusion that must be true must be a valid argument—no matter what the premises are! (If this confuses you, look back at the truth table for the conditional. Our definition of valid includes the conditional: if the premises are true, then the conclusion is true. Suppose now our conclusion must be true. Any conditional with a true consequent is true. So the definition of "valid" must be true of any argument with a tautology as a conclusion.) And, given that, it would seem that it is irrelevant whether we have any premises at all, since any will do. This suggests that there can be valid arguments with no premises.

Conditional derivation lets us actually construct such arguments. First, we will draw our Fitch bar for our main argument to indicate that we have no premises. Then, we will construct a conditional derivation. It will start like this:

1. $(P \rightarrow Q)$	assumption for conditional derivation

But what now? Well, we have assumed the antecedent of our sentence, and we should strive now to show the consequent. But note that the consequent is a conditional. So, we will again do a conditional derivation.

1. $(P \rightarrow Q)$	assumption for conditional derivation
2. $\neg Q$	assumption for conditional derivation
3. $\neg P$	modus tollens, 1, 2
4. $(\neg Q \rightarrow \neg P)$	conditional derivation, 2-3
5. $((P \rightarrow Q) \rightarrow (\neg Q \rightarrow \neg P))$	conditional derivation, 1-4

This is a proof, without premises, of $\mathbf{((P \rightarrow Q) \rightarrow (\neg Q \rightarrow \neg P))}$. The top of the proof shows that we have no premises. Our conclusion is a conditional, so, on line 1, we assumed the antecedent of the conditional. We now have to show the consequent of the conditional; but the consequent of the conditional is also a conditional, so we assumed its antecedent on line 2. Line 4 is the result of the conditional derivation from lines 2 to 3. Lines 1 through 4 tell us that if $\mathbf{(P \rightarrow Q)}$ is true, then $\mathbf{(\neg Q \rightarrow \neg P)}$ is true. And that is what we conclude on line 5.

We call a sentence that can be proved without premises a "theorem". Theorems are special because they reveal the things that follow from logic alone. It is a very great benefit of our propositional logic that all the theorems are tautologies. It is an equally great benefit of our propositional logic that all the tautologies are theorems. Nonetheless, these concepts are different. “Tautology” refers to a semantic concept: a tautology is a sentence that must be true. “Theorem” refers to a concept of syntax and derivation: a theorem is a sentence that can be derived without premises.

> Theorem: a sentence that can be proved without premises.
>
> Tautology: a sentence of the propositional logic that must be true.

6.5 Problems

1. Prove the following arguments are valid. This will require conditional derivation.

    a. Premise: $(\neg Q \rightarrow \neg P)$. Conclusion: $(P \rightarrow Q)$.
    b. Premise: $(P \rightarrow Q)$, $(Q \rightarrow R)$. Conclusion: $(P \rightarrow R)$.
    c. Premises: $(P \rightarrow (Q \rightarrow R))$, Q. Conclusion: $(P \rightarrow R)$.
    d. Premise: $(P \rightarrow Q)$, $(S \rightarrow R)$. Conclusion: $((\neg Q \wedge \neg R) \rightarrow (\neg P \wedge \neg S))$.
    e. Premise: $(P \rightarrow Q)$. Conclusion: $((P \wedge R) \rightarrow Q)$.
    f. Premise: $((R \wedge Q) \rightarrow S)$, $(\neg P \rightarrow (R \wedge Q))$. Conclusion: $(\neg S \rightarrow P)$.
    g. Premise: $(P \rightarrow \neg Q)$. Conclusion: $(Q \rightarrow \neg P)$.
    h. Premises: $(P \rightarrow Q)$, $(P \rightarrow R)$. Conclusion:$(P \rightarrow (Q \wedge R))$.
    i. Premises: $(P \rightarrow (Q \wedge S))$, $(Q \rightarrow R)$, $(S \rightarrow T)$. Conclusion: $(P \rightarrow (R \wedge T))$.
    j. Premises: $(P \rightarrow (Q \rightarrow R))$, $(P \rightarrow (S \rightarrow T))$,$(Q \wedge S)$. Conclusion: $(P \rightarrow (R \wedge T))$.

2. Prove the following theorems.

    a. $(P \rightarrow P)$.
    b. $((P \rightarrow Q) \rightarrow ((R \rightarrow P) \rightarrow (R \rightarrow Q)))$.
    c. $((P \rightarrow (Q \rightarrow R)) \rightarrow ((P \rightarrow Q) \rightarrow (P \rightarrow R))$.
    d. $((\neg P \rightarrow Q) \rightarrow (\neg Q \rightarrow P))$.
    e. $(((P \rightarrow Q) \wedge (P \rightarrow R)) \rightarrow (P \rightarrow (Q \wedge R)))$.

3. Make a truth table for each of the following complex sentences, in order to see when it is true or false. Identify which are tautologies. Prove the tau-

tologies.

a. $((P \rightarrow Q) \rightarrow Q)$.
b. $(P \rightarrow (Q \rightarrow Q))$.
c. $((P \rightarrow Q) \rightarrow P)$.
d. $(P \rightarrow (Q \rightarrow P))$.
e. $(P \rightarrow (P \rightarrow Q))$.
f. $((P \rightarrow P) \rightarrow Q))$.
g. $(P \rightarrow \neg P)$.
h. $(P \rightarrow \neg\neg P)$.
i. $((P \rightarrow Q) \rightarrow (P \wedge Q))$.
j. $((P \wedge Q) \rightarrow (P \rightarrow Q))$.

4. In normal colloquial English, write your own valid argument with at least two premises and with a conclusion that is a conditional. Your argument should just be a paragraph (not an ordered list of sentences or anything else that looks like formal logic). Translate it into propositional logic and prove it is valid.

5. Translate the following passage into our propositional logic. Prove the argument is valid.

> Either Beneke or Mill is the culprit who burned the Logician's Club. Also, if Beneke did it, then he bought the flares. But if Beneke bought the flares, he was at the Mariner's Shop yesterday. Thus, if Beneke was not at the Mariner's Shop yesterday, Mill did it.

---

[8] Hobbes (1886: 64).

7. “Or”

7.1 A historical example: The Euthryphro argument

The philosopher Plato (who lived from approximately 427 BC to 347 BC) wrote a series of great philosophical texts. Plato was the first philosopher to deploy argument in a vigorous and consistent way, and in so doing he showed how philosophy takes logic as its essential method. We think of Plato as the principal founder of Western philosophy. The American philosopher Alfred Whitehead (1861-1947) in fact once famously quipped that philosophy is a “series of footnotes to Plato”.

Plato’s teacher was Socrates (c. 469-399 B.C.), a gadfly of ancient Athens who made many enemies by showing people how little they knew. Socrates did not write anything, but most of Plato’s writings are dialogues, which are like small plays, in which Socrates is the protagonist of the philosophical drama that ensues. Several of the dialogues are named after the person who will be seen arguing with Socrates. In the dialogue *Euthyphro,* Socrates is standing in line, awaiting his trial. He has been accused of corrupting the youth of Athens. A trial in ancient Athens was essentially a debate before the assembled citizen men of the city. Before Socrates in line is a young man, Euthyphro. Socrates asks Euthyphro what his business is that day, and Euthyphro proudly proclaims he is there to charge his own father with murder. Socrates is shocked. In ancient Athens, respect for one’s father was highly valued and expected. Socrates, with characteristic sarcasm, tells Euthyphro that he must be very wise to be so confident. Here are two profound and conflicting duties: to respect one’s father, and to punish murder. Euthyphro seems to find it very easy to decide which is the greater duty. Euthyphro is not bothered. To him, these ethical matters are simple: one should be pious. When Socrates demands a definition of piety that applies to all pious acts, Euthyphro says,

> Piety is that which is loved by the gods and impiety is that which is not loved by them.

Socrates observes that this is ambiguous. It could mean, an act is good because the gods love that act. Or it could mean, the gods love an act because it is good. We have, then, an "or" statement, which logicians call a "disjunction":

> Either an act is good because the gods love that act, or the gods love an act because it is good.

Might the former be true? This view—that an act is good because the gods love it—is now called "divine command theory", and theists have disagreed since Socrates's time about whether it is true. But, Socrates finds it absurd. For, if tomorrow the gods love, say, murder, then, tomorrow murder would be good.

Euthyphro comes to agree that it cannot be that an act is good because the gods love that act. Our argument so far has this form:

> Either an act is good because the gods love that act, or the gods love an act because it is good.
>
> It is not the case that an act is good because the gods love it.

Socrates concludes that the gods love an act because it is good.

> Either an act is good because the gods love that act, or the gods love an act because it is good.
>
> It is not the case that an act is good because the gods love it.
>
> _____
>
> The gods love an act because it is good.

This argument is one of the most important arguments in philosophy. Most philosophers consider some version of this argument both valid and sound. Some who disagree with it bite the bullet and claim that if tomorrow God (most theistic philosophers alive today are monotheists) loved puppy torture, adultery, random acts of cruelty, pollution, and lying, these would all be good things. (If you are inclined to say, "That is not fair, God would never love those things", then you have already agreed with Socrates. For, the reason you believe that God would never love these kinds of acts is because these kinds of acts are bad. But then, being bad or good is something independent of the love of God.) But most philosophers agree with Socrates: they find it absurd to believe that random acts of cruelty and other such acts could be good. There is something inherently bad to these acts, they believe. The importance of the Euthyphro argument is not that it helps illustrate that divine command theory is an enormously strange and costly position to hold (though that is an important outcome), but rather that the argument shows ethics can be studied independently of theology. For, if there is something about acts that makes them good or bad independently of a god's will, then we do not have to study a god's will to study what makes those acts good or bad.

Of course, many philosophers are atheists so they already believed this, but for most of philosophy's history, one was obliged to be a theist. Even today, lay people tend to think of ethics as an extension of religion. Philosophers believe instead that ethics is its own field of study. The Euthyphro argument explains why, even if you are a theist, you can study ethics independently of studying theology.

But is Socrates's argument valid? Is it sound?

7.2 The disjunction

We want to extend our language so that it can represent sentences that contain an "or". Sentences like

> Tom will go to Berlin or Paris.
>
> We have coffee or tea.
>
> This web page contains the phrase "Mark Twain" or "Samuel Clemens."

Logicians call these kinds of sentences "disjunctions". Each of the two parts of a disjunction is called a "disjunct". The idea is that these are really equivalent to the following sentences:

> Tom will go to Berlin or Tom will go to Paris.
>
> We have coffee or we have tea.
>
> This web page contains the phrase "Mark Twain" or this web page contains the phrase "Samuel Clemens."

We can, therefore, see that (at least in many sentences) the "or" operates as a connective between two sentences.

It is traditional to use the symbol "**v**" for "or". This comes from the Latin "vel," meaning (in some contexts) *or.*

The syntax for the disjunction is very basic. If **Φ** and **Ψ** are sentences, then

(Φ v Ψ)

is a sentence.

The semantics is a little more controversial. This much of the defining truth table, most people find obvious:

Φ	Ψ	(ΦvΨ)
T	*T*	
T	*F*	*T*
F	*T*	*T*
F	*F*	*F*

Consider: if I promise that I will bring you roses or lilacs, then it seems that I told the truth either if I have brought you roses but not lilacs, or if I brought you lilacs but not roses. Similarly, the last row should be intuitive, also. If I promise I will bring you roses or lilacs, and I bring you nothing, then I spoke falsely.

What about the first row? Many people who are not logicians want it to be the case that we define this condition as false. The resulting meaning would correspond to what is sometimes called the "exclusive 'or'". Logicians disagree. They favor the definition where a disjunction is true if its two parts are true; this is sometimes called the "inclusive 'or'". Of course, all that matters is that we pick a definition and stick with it, but we can offer some reasons why the "inclusive 'or'", as we call it, is more general than the "exclusive 'or'".

Consider the first two sentences above. It seems that the first sentence—"Tom will go to Berlin or Paris"—should be true if Tom goes to both. Or consider the second sentence, "We have coffee or tea." In most restaurants, this means they have both coffee and they have tea, but they expect that you will order only one of these. After all, it would be strange to be told that they have coffee or tea, and then be told that it is false that they have both coffee and tea. Or, similarly, suppose the waiter said, “We have coffee or tea”, and then you said “I’ll have both”, and the waiter replied “We don’t have both”. This would seem strange. But if you find it strange, then you implicitly agree that the disjunction should be interpreted as the inclusive “or”.

Examples like these suggest to logicians that the inclusive “or” (where the first row of the table is true) is the default case, and that the context of our speech tells us when not both disjuncts are true. For example, when a restaurant has a fixed price menu—where you pay one fee and then get either steak or lobster—it is understood by the context that this means you can have one or the other but not both. But that is not logic, that is social custom. One must know about restaurants to determine this.

Thus, it is customary to define the semantics of the disjunction as

Φ	**Ψ**	**(ΦvΨ)**
T	*T*	*T*
T	*F*	*T*
F	*T*	*T*
F	*F*	*F*

We haven’t lost the ability to express the exclusive “or”. We can say, “one or the other but not both”, which is expressed by the formula “**((Φ v Ψ) ^ ¬(Φ ^ Ψ))**”. To check, we can make the truth table for this complex expression:

Φ	**Ψ**	**(Φ ^ Ψ)**	**(Φ v Ψ)**	**¬(Φ ^ Ψ)**	**((Φ v Ψ) ^ ¬(Φ ^ Ψ))**
T	*T*	*T*	*T*	*F*	*F*
T	*F*	*F*	*T*	*T*	*T*
F	*T*	*F*	*T*	*T*	*T*
F	*F*	*F*	*F*	*T*	*F*

Note that this formula is equivalent to the exclusive “or” (it is true when **Φ** is true or **Ψ** is true, but not when both are true or both are false). So, if we need to say something like the exclusive “or”, we can do so.

7.3 Alternative forms

There do not seem to be many alternative expressions in English equivalent to the "or". We have

P or **Q**

Either **P** or **Q**

These are both expressed in our logic with **(P v Q)**.

One expression that does arise in English is "neither…nor…". This expression seems best captured by simply making it into "not either… or…". Let's test this proposal. Consider the sentence

Neither Smith nor Jones will go to London.

This sentence expresses the idea that Smith will not go to London, and that Jones will not go to London. So, it would surely be a mistake to express it as

Either Smith will not go to London or Jones will not go to London.

Why? Because this latter sentence would be true if one of them went to London and one of them did not. Consider the truth table for this expression to see this. Use the following translation key.

P: Smith will go to London.

Q: Jones will go to London.

Then suppose we did (wrongly) translate "Neither Smith nor Jones will go to London" with

(¬P v ¬Q)

Here is the truth table for this expression.

P	Q	¬Q	¬P	(¬Pv¬Q)
T	*T*	*F*	*F*	*F*
T	*F*	*T*	*F*	*T*
F	*T*	*F*	*T*	*T*
F	*F*	*T*	*T*	*T*

Note that this sentence is true if **P** is true and **Q** is false, or if **Q** is true and **P** is false. In other words, it is true if one of the two goes to London. That's not what we mean in English by that sentence claiming that neither of them will go to London.

The better translation is **¬(PvQ)**.

P	Q	(PvQ)	¬(PvQ)
T	*T*	*T*	*F*
T	*F*	*T*	*F*
F	*T*	*T*	*F*
F	*F*	*F*	*T*

This captures the idea well: it is only true if each does not go to London. So, we can simply translate "neither…nor…" as "It is not the case that either… or…".

7.4 Reasoning with disjunctions

How shall we reason with the disjunction? Looking at the truth table that defines the disjunction, we find that we do not know much if we are told that, say, **(P v Q)**. **P** could be true, or it could be false. The same is so for **Q**. All we know is that they cannot both be false.

This does suggest a reasonable and useful kind of inference rule. If we have a disjunction, and we discover that half of it is false, then we know that the other half must be true. This is true for either disjunct. This means we have two rules, but we can group together both rules with a single name and treat them as one rule:

(Φ v Ψ)

¬Φ

_____

Ψ

and

(Φ v Ψ)

¬Ψ

_____

Φ

This rule is traditionally called “modus tollendo ponens”.

What if we are required to show a disjunction? One insight we can use is that if some sentence is true, then any disjunction that contains it is true. This is so whether the sentence makes up the first or second disjunct. Again, then, we would have two rules, which we can group together under one name:

Φ

_____

(Φ v Ψ)

and

Ψ

_____

(Φ v Ψ)

This rule is often called “addition”.

The addition rule often confuses students. It seems to be a cheat, as if we are getting away with something for free. But a moment of reflection will help clarify that just the opposite is true. We lose information when we use the addition rule.

If you ask me where John is, and I say, “John is in New York”, I told you more than if I answered you, “John is either in New York or in New Jersey”. Just so, when we go from some sentence **P** to **(PvQ)**, we did not get something for free.

This rule does have the seemingly odd consequence that from, say, 2+2=4 you can derive that either 2+2=4 or 7=0. But that only seems odd because in normal speech, we have a number of implicit rules. The philosopher Paul Grice (1913-1988) described some of these rules, and we sometimes call the rules he described “Grice’s Maxims”.[9] He observed that in conversation we expect people to give all the information required but not more; to try to be truthful; to say things that are relevant; and to be clear and brief and orderly. So, in normal English conversations, if someone says, “Tom is in New York or New Jersey,” they would be breaking the rule to give enough information, and to say what is relevant, if they knew that Tom was in New York. This also means that we expect people to use a disjunction when they have reason to believe that either or both disjuncts could be true. But our logical language is designed only to be precise, and we have been making the language precise by specifying when a sentence is true or false, and by specifying the relations between sentences in terms of their truth values. We are thus not representing, and not putting into our language, Grice’s maxims of conversation. It remains true that if you knew Tom is in New York, but answered my question “Where is Tom?” by saying “Tom is in New York or New Jersey”, then you have wasted my time. But you did not say something false.

We are now in a position to test Socrates’s argument. Using the following translation key, we can translate the argument into symbolic form.

P: An act is good because the gods love that act.

Q: The gods love an act because it is good.

Euthyphro had argued

1. $(P \vee Q)$	premise

Socrates had got Euthryphro to admit that

1. $(P \vee Q)$	premise
2. $\neg P$	premise

And so we have a simple direct derivation:

1. $(P \vee Q)$	premise
2. $\neg P$	premise
3. Q	modus tollendo ponens, 1, 2

Socrates's argument is valid. I will leave it up to you to determine whether Socrates's argument is sound.

Another example might be helpful. Here is an argument in our logical language.

(P v Q)

¬P

(¬P → (Q → R))

_____

(R v S)

This will make use of the addition rule, and so is useful to illustrating that rule's application. Here is one possible proof.

1. $(P \vee Q)$	premise
2. $\neg P$	premise
3. $(\neg P \rightarrow (Q \rightarrow R))$	premise
4. Q	modus tollendo ponens, 1, 2
5. $(Q \rightarrow R)$	modus ponens, 3, 2
6. R	modus ponens, 5, 4
7. $(R \vee S)$	addition, 6

7.5 Alternative symbolizations of disjunction

We are fortunate that there have been no popular alternatives to the use of "**v**" as a symbol for disjunction. Perhaps the second most widely used alternative symbol was "**||**", such that **(P v Q)** would be symbolized:

(P || Q)

7.6 Problems

1. Prove the following using a derivation.

    a. Premises: **(PvQ)**, **(Q → S)**, **(¬S^T)**. Conclusion: **(T^P)**.
    b. Premises: **((P → ¬Q) ^ (R → S))**, **(Q v R)**. Conclusion: **(P → S)**.
    c. Premises: **((P^Q) v R)**, **((P^Q) → S)**, **¬S**. Conclusion: **R**.
    d. Premises: **(RvS)**, **((S → T) ^ V)**, **¬T**, **((R^V) → P)**. Conclusion: **(PvQ)**.
    e. Premises: **((P → Q) v (¬R → S))**, **((P → Q) → T)**, **(¬T ^ ¬S)**. Conclusion: **(R v V)**.
    f. Premises: **(P v S)**, **(T → ¬S)**, **T**. Conclusion: **((P v Q) v R)**.
    g. Conclusion: **(P → (PvQ))**.
    h. Conclusion: **((PvQ) → (¬P → Q))**.
    i. Conclusion: **((PvQ) → (¬Q → P))**.
    j. Conclusion: **(((PvQ) ^ (¬Q v ¬R)) → (R → P))**.

2. Consider the following four cards in figure 7.1. Each card has a letter on one side, and a shape on the other side.

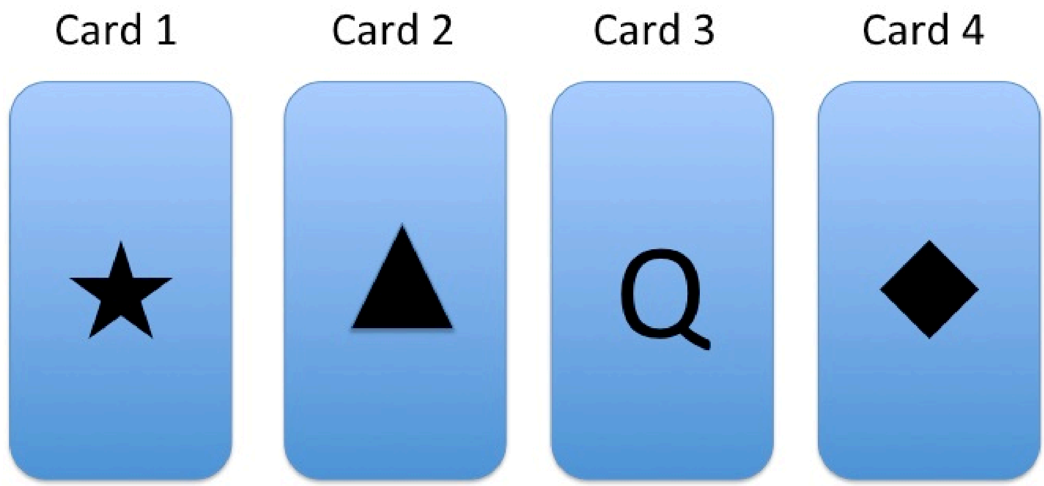

Figure 7.1

For each of the following claims, determine (1) the minimum number of cards you must turn over to check the claim, and (2) what those cards are, in order to determine if the claim is true of all four cards.

a. If there is a P or Q on the letter side of the card, then there is a diamond on the shape side of the card.
b. If there is a Q on the letter side of the card, then there is either a diamond or a star on the shape side of the card.

3. In normal colloquial English, write your own valid argument with at least two premises, at least one of which is a disjunction. Your argument should just be a paragraph (not an ordered list of sentences or anything else that looks like formal logic). Translate it into propositional logic and prove it is valid.

4. Translate the following passage into our propositional logic. Prove the argument is valid.

   Either Dr. Kronecker or Bishop Berkeley killed Colonel Cardinality. If Dr. Kronecker killed Colonel Cardinality, then Dr. Kronecker was in the kitchen. If Bishop Berkeley killed Colonel Cardinality, then he was in the drawing room. If Bishop Berkeley was in the drawing room, then he was wearing boots. But Bishop Berkeley was not wearing boots. So, Dr. Kronecker killed the Colonel.

5. Translate the following passage into our propositional logic. Prove the argument is valid.

> Either Wittgenstein or Meinong stole the diamonds. If Meinong stole the diamonds, then he was in the billiards room. But if Meinong was in the library, then he was not in the billiards room. Therefore, if Meinong was in the library, Wittgenstein stole the diamonds.

---

[9] Grice (1975).

8. Reductio ad Absurdum

8.1 A historical example

In his book, *The Two New Sciences*,[10] Galileo Galilea (1564-1642) gives several arguments meant to demonstrate that there can be no such thing as actual infinities or actual infinitesimals. One of his arguments can be reconstructed in the following way. Galileo proposes that we take as a premise that there is an actual infinity of natural numbers (the natural numbers are the positive whole numbers from 1 on):

{1, 2, 3, 4, 5, 6, 7, ….}

He also proposes that we take as a premise that there is an actual infinity of the squares of the natural numbers.

{1, 4, 9, 16, 25, 36, 49, ….}

Now, Galileo reasons, note that these two groups (today we would call them "sets") have the same size. We can see this because we can see that there is a one-to-one correspondence between the two groups.

{1,	2,	3,	4,	5,	6,	7, ….}
↕	↕	↕	↕	↕	↕	↕
{1,	4,	9,	16,	25,	36,	49, …}

If we can associate every natural number with one and only one square number, and if we can associate every square number with one and only one natural number, then these sets must be the same size.

But wait a moment, Galileo says. There are obviously very many more natural numbers than there are square numbers. That is, every square number is in the list of natural numbers, but many of the natural numbers are not in the list of square numbers. The following numbers are all in the list of natural numbers but not in the list of square numbers.

{2, 3, 5, 6, 7, 8, 10, ...}

So, Galileo reasons, if there are many numbers in the group of natural numbers that are not in the group of the square numbers, and if there are no numbers in the group of the square numbers that are not in the naturals numbers, then the natural numbers is bigger than the square numbers. And if the group of the natural numbers is bigger than the group of the square numbers, then the natural numbers and the square numbers are not the same size.

We have reached two conclusions: the set of the natural numbers and the set of the square numbers are the same size; and, the set of the natural numbers and the set of the square numbers are not the same size. That's contradictory.

Galileo argues that the reason we reached a contradiction is because we assumed that there are actual infinities. He concludes, therefore, that there are no actual infinities.

8.2 Indirect proofs

Our logic is not yet strong enough to prove some valid arguments. Consider the following argument as an example.

(P→(QvR))

¬Q

¬R

_____

¬P

This argument looks valid. By the first premise we know: if **P** were true, then so would **(Q v R)** be true. But then either **Q** or **R** or both would be true. And by the second and third premises we know: **Q** is false and **R** is false. So it cannot be that **(Q v R)** is true, and so it cannot be that **P** is true.

We can check the argument using a truth table. Our table will be complex because one of our premise is complex.

				premise	premise	premise	conclusion
P	Q	R	(QvR)	(P→(QvR))	¬Q	¬R	¬P
T	*T*	*T*	*T*	*T*	*F*	*F*	*F*
T	*T*	*F*	*T*	*T*	*F*	*T*	*F*
T	*F*	*T*	*T*	*T*	*T*	*F*	*F*
T	*F*	*F*	*F*	*F*	*T*	*T*	*F*
F	*T*	*T*	*T*	*T*	*F*	*F*	*T*
F	*T*	*F*	*T*	*T*	*F*	*T*	*T*
F	*F*	*T*	*T*	*T*	*T*	*F*	*T*
F	*F*	*F*	*F*	*T*	*T*	*T*	*T*

In any kind of situation in which all the premises are true, the conclusion is true. That is: the premises are all true only in the last row. For that row, the conclusion is also true. So, this is a valid argument.

But take a minute and try to prove this argument. We begin with

1. $(\mathsf{P} \rightarrow (\mathsf{Q} \vee \mathsf{R}))$	premise
2. $\neg\mathsf{Q}$	premise
3. $\neg\mathsf{R}$	premise

And now we are stopped. We cannot apply any of our rules. Here is a valid argument that we have not made our reasoning system strong enough to prove.

There are several ways to rectify this problem and to make our reasoning system strong enough. One of the oldest solutions is to introduce a new proof method, traditionally called "reductio ad absurdum", which means *a reduction to absurdity*. This method is also often called an "indirect proof" or "indirect derivation".

The idea is that we assume the denial of our conclusion, and then show that a contradiction results. A contradiction is shown when we prove some sentence **Ψ**, and its negation ¬**Ψ**. This can be any sentence. The point is that, given the principle of bivalence, we must have proven something false. For if **Ψ** is true, then ¬**Ψ** is false;

and if ¬Ψ is true, then Ψ is false. We don't need to know which is false (Ψ or ¬Ψ); it is enough to know that one of them must be.

Remember that we have built our logical system so that it cannot produce a falsehood from true statements. The source of the falsehood that we produce in the indirect derivation must, therefore, be some falsehood that we added to our argument. And what we added to our argument is the denial of the conclusion. Thus, the conclusion must be true.

The shape of the argument is like this:

$$\begin{array}{l|l}
 & \quad \neg\phi \\
 & \quad \vdots \\
 & \quad \psi \\
 & \quad \neg\psi \\
\triangleright & \phi
\end{array}$$

Traditionally, the assumption for indirect derivation has also been commonly called "the assumption for reductio".

As a concrete example, we can prove our perplexing case.

1. $(\mathsf{P} \rightarrow (\mathsf{Q} \vee \mathsf{R}))$	premise
2. $\neg\mathsf{Q}$	premise
3. $\neg\mathsf{R}$	premise
4. $\neg\neg\mathsf{P}$	assumption for indirect derivation
5. P	double negation, 4
6. $(\mathsf{Q} \vee \mathsf{R})$	modus ponens, 1, 5
7. R	modus tollendo ponens, 6, 2
8. $\neg\mathsf{R}$	repeat, 3
9. $\neg\mathsf{P}$	indirect derivation, 4-8

We assumed the denial of our conclusion on line 4. The conclusion we believed was correct was ¬P, and the denial of this is ¬¬P. In line 7, we proved R. Technically, we are done at that point, but we would like to be kind to anyone trying to understand our proof, so we repeat line 3 so that the sentences R and ¬R are side by side, and it is very easy to see that something has gone wrong. That is, if we have proven both **R** and **¬R**, then we have proven something false.

Our reasoning now goes like this. What went wrong? Line 8 is a correct use of repetition; line 7 comes from a correct use of modus tollendo ponens; line 6 from a correct use of modus ponens; line 5 from a correct use of double negation. So, we did not make a mistake in our reasoning. We used lines 1, 2, and 3, but those are premises that we agreed to assume are correct. This leaves line 4. That must be the source of my contradiction. It must be false. If line 4 is false, then **¬P** is true.

Some people consider indirect proofs less strong than direct proofs. There are many, and complex, reasons for this. But, for our propositional logic, none of these reasons apply. This is because it is possible to prove that our propositional logic is consistent. This means, it is possible to prove that our propositional logic cannot prove a falsehood unless one first introduces a falsehood into the system. (It is generally not possible to prove that more powerful and advanced logical or mathematical systems are consistent, from inside those systems; for example, one cannot prove in arithmetic that arithmetic is consistent.) Given that we can be certain of the consistency of the propositional logic, we can be certain that in our propositional logic an indirect proof is a good form of reasoning. We know that if we prove a falsehood, we must have put a falsehood in; and if we are confident about all the other assumptions (that is, the premises) of our proof except for the assumption for indirect derivation, then we can be confident that this assumption for indirect derivation must be the source of the falsehood.

A note about terminology is required here. The word "contradiction" gets used ambiguously in most logic discussions. It can mean a situation like we see above, where two sentences are asserted, and these sentences cannot both be true. Or it can mean a single sentence that cannot be true. An example of such a sentence is **(P^¬P)**. The truth table for this sentence is:

P	¬P	(P ^ ¬P)
T	*F*	*F*
F	*T*	*F*

Thus, this kind of sentence can never be true, regardless of the meaning of **P**.

To avoid ambiguity, in this text, we will always call a single sentence that cannot be true a “contradictory sentence”. Thus, **(P^¬P)** is a contradictory sentence. Situations where two sentences are asserted that cannot both be true will be called a “contradiction”.

8.3 Our example, and other examples

We can reconstruct a version of Galileo’s argument now. We will use the following key.

> **P**: There are actual infinities (including the natural numbers and the square numbers).
>
> **Q**: There is a one-to-one correspondence between the natural numbers and the square numbers.
>
> **R**: The size of the set of the natural numbers and the size of the set of the square numbers are the same.
>
> **S**: All the square numbers are natural numbers.
>
> **T**: Some of the natural numbers are not square numbers.
>
> **U**: There are more natural numbers than square numbers.

With this key, the argument will be translated:

> **(P→Q)**
>
> **(Q→R)**
>
> **(P→(S^T))**
>
> **((S^T)→U)**
>
> **(U→¬R)**
>
> ______
>
> **¬P**

And we can prove this is a valid argument by using indirect derivation:

1. $(P \rightarrow Q)$	premise
2. $(Q \rightarrow R)$	premise
3. $(P \rightarrow (S \wedge T))$	premise
4. $((S \wedge T) \rightarrow U)$	premise
5. $(U \rightarrow \neg R)$	premise
6. $\neg\neg P$	assumption for indirect derivation
7. P	double negation, 6
8. Q	modus ponens, 1, 7
9. R	modus ponens, 2, 8
10. $(S \wedge T)$	modus ponens, 3, 7
11. U	modus ponens, 4, 10
12. $\neg R$	modus ponens, 5, 11
13. R	repeat, 9
14. $\neg P$	indirect derivation, 6-13

On line 6, we assumed **¬¬P** because Galileo believed that **¬P** and aimed to prove that **¬P**. That is, he believed that there are no actual infinities, and so assumed that it was false to believe that it is not the case that there are no actual infinities. This falsehood will lead to other falsehoods, exposing itself.

For those who are interested: Galileo concluded that there are no actual infinities but there are potential infinities. Thus, he reasoned, it is not the case that all the natural numbers exist (in some sense of "exist"), but it is true that you could count natural numbers forever. Many philosophers before and after Galileo held this view; it is similar to a view held by Aristotle, who was an important logician and philosopher writing nearly two thousand years before Galileo.

Note that in an argument like this, you could reason that not the assumption for indirect derivation, but rather one of the premises was the source of the contradiction. Today, most mathematicians believe this about Galileo's argument. A logician and mathematician named Georg Cantor (1845-1918), the inventor of set theory, argued that infinite sets can have proper subsets of the same size. That is, Cantor denied premise 4 above: even though all the square numbers are natural numbers, and not all natural numbers are square numbers, it is not the case that these two sets are of different size. Cantor accepted however premise 2 above, and, therefore, believed that the size of the set of natural numbers and the size of the set of square numbers is the same. Today, using Cantor's reasoning, mathematicians and logicians study infinity, and have developed a large body of knowledge about the nature of infinity. If this interests you, see section 17.5.

Let us consider another example to illustrate indirect derivation. A very useful set of theorems are today called “De Morgan’s Theorems”, after the logician Augustus De Morgan (1806–1871). We cannot state these fully until chapter 9, but we can state their equivalent in English: DeMorgan observed that **¬(PvQ)** and **(¬P^¬Q)** are equivalent, and also that **¬(P^Q)** and **(¬Pv¬Q)** are equivalent. Given this, it should be a theorem of our language that **(¬(PvQ)→(¬P^¬Q))**. Let’s prove this.

The whole formula is a conditional, so we will use a conditional derivation. Our proof must thus begin:

1. $\neg(P \vee Q)$	assumption for conditional derivation

To complete the conditional derivation, we must prove **(¬P^¬Q)**. This is a conjunction, and our rule for showing conjunctions is adjunction. Since using this rule might be our best way to show **(¬P^¬Q)**, we can aim to show **¬P** and then show **¬Q**, and then perform adjunction. But, we obviously have very little to work with—just line 1, which is a negation. In such a case, it is typically wise to attempt an indirect proof. Start with an indirect proof of **¬P**.

1. $\neg(P \vee Q)$	assumption for conditional derivation
2. $\neg\neg P$	assumption for indirect derivation
3. P	double negation, 2

We now need to find a contradiction—any contradiction. But there is an obvious one already. Line 1 says that neither **P** nor **Q** is true. But line 3 says that **P** is true. We must make this contradiction explicit by finding a formula and its denial. We can do this using addition.

1. $\neg(P \lor Q)$	assumption for conditional derivation
2. $\neg\neg P$	assumption for indirect derivation
3. P	double negation, 2
4. $(P \lor Q)$	addition, 3
5. $\neg(P \lor Q)$	repeat, 1
6. $\neg P$	indirect derivation 2-5

To complete the proof, we will use this strategy again.

1. $\neg(P \lor Q)$	assumption for conditional derivation
2. $\neg\neg P$	assumption for indirect derivation
3. P	double negation, 2
4. $(P \lor Q)$	addition, 3
5. $\neg(P \lor Q)$	repeat, 1
6. $\neg P$	indirect derivation 2-5
7. $\neg\neg Q$	assumption for indirect derivation
8. Q	double negation, 7
9. $(P \lor Q)$	addition, 8
10. $\neg(P \lor Q)$	repeat, 1
11. $\neg Q$	indirect derivation 7-10
12. $(\neg P \land \neg Q)$	adjunction, 6, 11
13. $(\neg(P \lor Q) \rightarrow (\neg P \land \neg Q))$	conditional derivation 1-12

We will prove De Morgan's theorems as problems for chapter 9.

Here is a general rule of thumb for doing proofs: When proving a conditional, always do conditional derivation; otherwise, try direct derivation; if that fails, then, try indirect derivation.

8.4 Problems

1. Complete the following proofs. Each will require an indirect derivation. The last two are challenging.

a. Premises: $(P \rightarrow R)$, $(Q \rightarrow R)$, (PvQ). Conclusion: R.
b. Premises: $((PvQ) \rightarrow R)$, $\neg R$. Conclusion: $\neg P$.
c. Premise: $(\neg P \wedge \neg Q)$. Conclusion: $\neg(PvQ)$.
d. Premise: $\neg(P \wedge Q)$. Conclusion: $(\neg Pv\neg Q)$.
e. Premise: $(\neg Pv\neg Q)$. Conclusion: $\neg(P \wedge Q)$.
f. Premise: $(P \rightarrow R)$, $(Q \rightarrow S)$, $\neg(R \wedge S)$. Conclusion: $\neg(P \wedge Q)$.
g. Premise: $\neg R$, $((P \rightarrow R) \vee (Q \rightarrow R))$. Conclusion: $(\neg P \vee \neg Q)$.
h. Premise: $\neg(R \vee S)$, $(P \rightarrow R)$, $(Q \rightarrow S)$. Conclusion: $\neg(P \vee Q)$.

2. Prove the following are theorems.

   a. $\neg(P \wedge \neg P)$.
   b. $(\neg P \rightarrow \neg(P \wedge Q))$.
   c. $((P \wedge \neg Q) \rightarrow \neg(P \rightarrow Q))$.
   d. $((P \rightarrow Q) \rightarrow \neg(P \wedge \neg Q))$.
   e. $(\neg(P \vee Q) \rightarrow \neg P)$.
   f. $((\neg P \wedge \neg Q) \rightarrow \neg(P \vee Q))$.
   g. $((\neg P \vee \neg Q) \rightarrow \neg(P \wedge Q))$.
   h. $(\neg(P \wedge Q) \rightarrow (\neg P \vee \neg Q))$.
   i. $\neg((P \rightarrow \neg P) \wedge (\neg P \rightarrow P))$.
   j. $(P \vee \neg P)$.

3. In normal colloquial English, write your own valid argument with at least two premises. Your argument should just be a paragraph (not an ordered list of sentences or anything else that looks like formal logic). Translate it into propositional logic and prove it is valid using an indirect derivation.

4. Translate the following argument into English, and then prove it is valid using an indirect proof.

   > Either Beneke or Meinong conned Dodgson with marked cards. But either Beneke didn't do it or the marked cards are in the car. But also, if Meinong did it, the marked cards are in the car. So, regardless of whether Beneke or Meinong conned Dodgson, the cards are in the car.

[10] This translation of the title of Galileo's book has become the most common, although a more literal one would have been *Mathematical Discourses and Demonstrations*. Translations of the book include Drake (1974).

9. "… if and only if …", Using Theorems

9.1 A historical example

The philosopher David Hume (1711-1776) is remembered for being a brilliant skeptical empiricist. A person is a skeptic about a topic if that person both has very strict standards for what constitutes knowledge about that topic and also believes we cannot meet those strict standards. Empiricism is the view that we primarily gain knowledge through experience, particular experiences of our senses. In his book, *An Inquiry Concerning Human Understanding,* Hume lays out his principles for knowledge, and then advises us to clean up our libraries:

> When we run over libraries, persuaded of these principles, what havoc must we make? If we take in our hand any volume of divinity or school metaphysics, for instance, let us ask, Does it contain any abstract reasoning concerning quantity or number? No. Does it contain any experimental reasoning concerning matter of fact and existence? No. Commit it then to the flames, for it can contain nothing but sophistry and illusion.[11]

Hume felt that the only sources of knowledge were logical or mathematical reasoning (which he calls above "abstract reasoning concerning quantity or number") or sense experience ("experimental reasoning concerning matter of fact and existence"). Hume is led to argue that any claims not based upon one or the other method is worthless.

We can reconstruct Hume's argument in the following way. Suppose ***t*** is some topic about which we claim to have knowledge. Suppose that we did not get this knowledge from experience or logic. Written in English, we can reconstruct his argument in the following way:

> We have knowledge about ***t*** if and only if our claims about ***t*** are learned from experimental reasoning or from logic or mathematics.
>
> Our claims about ***t*** are not learned from experimental reasoning.

Our claims about ***t*** are not learned from logic or mathematics.

_____

We do not have knowledge about ***t.***

What does that phrase "if and only if" mean? Philosophers think that it, and several synonymous phrases, are used often in reasoning. Leaving "if and only" unexplained for now, we can use the following translation key to write up the argument in a mix of our propositional logic and English.

P: We have knowledge about ***t.***

Q: Our claims about ***t*** are learned from experimental reasoning.

R: Our claims about ***t*** are learned from logic or mathematics.

And so we have:

P if and only if **(QvR)**

¬Q

¬R

_____

¬P

Our task is to add to our logical language an equivalent to "if and only if". Then we can evaluate this reformulation of Hume's argument.

9.2 The biconditional

Before we introduce a symbol synonymous with "if and only if", and then lay out its syntax and semantics, we should start with an observation. A phrase like "**P** if and only if **Q**" appears to be an abbreviated way of saying "**P** if **Q** and **P** only if **Q**". Once we notice this, we do not have to try to discern the meaning of "if and only if" using our expert understanding of English. Instead, we can discern the meaning of "if and only if" using our already rigorous definitions of "if", "and", and "only if". Specifically, "**P** if **Q** and **P** only if **Q**" will be translated "**((Q→P)^(P→Q))**". (If this is unclear to you, go back and review section 2.2.) Now, let us make a truth table for this formula.

P	Q	(Q → P)	(P → Q)	((Q→P)^(P→Q))
T	*T*	*T*	*T*	*T*
T	*F*	*T*	*F*	*F*
F	*T*	*F*	*T*	*F*
F	*F*	*T*	*T*	*T*

We have settled the semantics for "if and only if". We can now introduce a new symbol for this expression. It is traditional to use the double arrow, "↔". We can now express the syntax and semantics of "↔".

If **Φ** and **Ψ** are sentences, then

(Φ↔Ψ)

is a sentence. This kind of sentence is typically called a "biconditional".

The semantics is given by the following truth table.

Φ	Ψ	(Φ↔Ψ)
T	*T*	*T*
T	*F*	*F*
F	*T*	*F*
F	*F*	*T*

One pleasing result of our account of the biconditional is that it allows us to succinctly explain the syntactic notion of logical equivalence. We say that two sentences **Φ** and **Ψ** are "equivalent" or "logically equivalent" if **(Φ↔Ψ)** is a theorem.

9.3 Alternative phrases

In English, it appears that there are several phrases that usually have the same meaning as the biconditional. Each of the following sentences would be translated as **$(P \leftrightarrow Q)$**.

P if and only if Q.

P just in case Q.

P is necessary and sufficient for Q.

P is equivalent to Q.

9.4 Reasoning with the biconditional

How can we reason using a biconditional? At first, it would seem to offer little guidance. If I know that **$(P \leftrightarrow Q)$**, I know that P and Q have the same truth value, but from that sentence alone I do not know if they are both true or both false. Nonetheless, we can take advantage of the semantics for the biconditional to observe that if we also know the truth value of one of the sentences constituting the biconditional, then we can derive the truth value of the other sentence. This suggests a straightforward set of rules. These will actually be four rules, but we will group them together under a single name, "equivalence":

$(\Phi \leftrightarrow \Psi)$

Φ

_____

Ψ

and

$(\Phi \leftrightarrow \Psi)$

Ψ

———

Φ

and

(Φ↔Ψ)

¬Φ

———

¬Ψ

and

(Φ↔Ψ)

¬Ψ

———

¬Φ

What if we instead are trying to show a biconditional? Here we can return to the insight that the biconditional **(Φ↔Ψ)** is equivalent to **((Φ→Ψ)^(Ψ→Φ))**. If we could prove both **(Φ→Ψ)** and **(Ψ→Φ)**, we will know that **(Φ↔Ψ)** must be true.

We can call this rule "bicondition". It has the following form:

(Φ→Ψ)

(Ψ→Φ)

———

(Φ↔Ψ)

This means that often when we aim to prove a biconditional, we will undertake two conditional derivations to derive two conditionals, and then use the bicondition rule. That is, many proofs of biconditionals have the following form:

$$
\begin{array}{l|l}
 & \quad | \ \phi \\
 & \quad | \ \vdots \\
 & \quad | \ \psi \\
\triangleright & (\phi \rightarrow \psi) \\
 & \quad | \ \psi \\
 & \quad | \ \vdots \\
 & \quad | \ \phi \\
\triangleright & (\psi \rightarrow \phi) \\
 & (\phi \leftrightarrow \psi)
\end{array}
$$

9.5 Returning to Hume

We can now see if we are able to prove Hume's argument. Given now the new biconditional symbol, we can begin a direct proof with our three premises.

$$
\begin{array}{|lr}
1.\ (\mathsf{P} \leftrightarrow (\mathsf{Q} \vee \mathsf{R})) & \text{premise} \\
2.\ \neg\mathsf{Q} & \text{premise} \\
3.\ \neg\mathsf{R} & \text{premise} \\
 & \\
\end{array}
$$

We have already observed that we think **(QvR)** is false because **¬Q** and **¬R.** So let's prove **¬(QvR).** This sentence cannot be proved directly, given the premises we have; and it cannot be proven with a conditional proof, since it is not a conditional. So let's try an indirect proof. We believe that **¬(QvR)** is true, so we'll assume the denial of this and show a contradiction.

1. $(P \leftrightarrow (Q \lor R))$	premise
2. $\neg Q$	premise
3. $\neg R$	premise
4. $\neg\neg(Q \lor R)$	assumption for indirect derivation
5. $(Q \lor R)$	double negation, 4
6. R	modus tollendo ponens, 5, 2
7. $\neg R$	repetition, 3
8. $\neg(Q \lor R)$	indirect proof, 4-7
9. $\neg P$	equivalence, 1, 8

Hume's argument, at least as we reconstructed it, is valid.

Is Hume's argument sound? Whether it is sound depends upon the first premise above (since the second and third premises are abstractions about some topic ***t***). Most specifically, it depends upon the claim that we have knowledge about something just in case we can show it with experiment or logic. Hume argues we should distrust—indeed, we should burn texts containing—claims that are not from experiment and observation, or from logic and math. But consider this claim: we have knowledge about a topic ***t*** if and only if our claims about ***t*** are learned from experiment or our claims about ***t*** are learned from logic or mathematics.

Did Hume discover this claim through experiments? Or did he discover it through logic? What fate would Hume's book suffer, if we took his advice?

9.6 Some examples

It can be helpful to prove some theorems that make use of the biconditional, in order to illustrate how we can reason with the biconditional.

Here is a useful principle. If two sentences have the same truth value as a third sentence, then they have the same truth value as each other. We state this as **$(((P \leftrightarrow Q) \wedge (R \leftrightarrow Q)) \rightarrow (P \leftrightarrow R))$**. To illustrate reasoning with the biconditional, let us prove this theorem.

This theorem is a conditional, so it will require a conditional derivation. The consequent of the conditional is a biconditional, so we will expect to need two conditional derivations, one to prove **(P→R)** and one to prove **(R→P)**. The proof will look like this. Study it closely.

1. $((P \leftrightarrow Q) \wedge (R \leftrightarrow Q))$	assumption for conditional derivation
2. $(P \leftrightarrow Q)$	simplification, 1
3. $(R \leftrightarrow Q)$	simplification, 1
4. P	assumption for conditional derivation
5. Q	equivalence, 2, 4
6. R	equivalence, 3, 5
7. $(P \rightarrow R)$	conditional derivation, 4-6
8. R	assumption for conditional derivation
9. Q	equivalence, 3, 8
10. P	equivalence, 2, 9
11. $(R \rightarrow P)$	conditional derivation, 8-10
12. $(P \leftrightarrow R)$	bincondition, 7, 11
13. $(((P \leftrightarrow Q) \wedge (R \leftrightarrow Q)) \rightarrow (P \leftrightarrow R))$	conditional derivation 1-12

We have mentioned before the principles that we associate with the mathematician Augustus De Morgan (1806-1871), and which today are called "De Morgan's Laws" or the "De Morgan Equivalences". These are the recognition that **¬(PvQ)** and **(¬P^¬Q)** are equivalent, and also that **¬(P^Q)** and **(¬Pv¬Q)** are equivalent. We can now express these with the biconditional. The following are theorems of our logic:

(¬(PvQ)↔(¬P^¬Q))

(¬(P^Q)↔(¬Pv¬Q))

We will prove the second of these theorems. This is perhaps the most difficult proof we have seen; it requires nested indirect proofs, and a fair amount of cleverness in finding what the relevant contradiction will be.

1. $\neg(P \wedge Q)$	assumption for conditional derivation
2. $\neg(\neg P \vee \neg Q)$	assumption for indirect derivation
3. $\neg P$	assumption for indirect derivation
4. $(\neg P \vee \neg Q)$	addition, 3
5. $\neg(\neg P \vee \neg Q)$	repeat, 2
6. P	indirect derivation, 3-5
7. $\neg Q$	assumption for indirect derivation
8. $(\neg P \vee \neg Q)$	addition, 7
9. $\neg(\neg P \vee \neg Q)$	repeat, 2
10. Q	indirect derivation, 6-9
11. $(P \wedge Q)$	adjunction, 6, 10
12. $\neg(P \wedge Q)$	repeat, 1
13. $(\neg P \vee \neg Q)$	indirect derivation, 2-12
14. $(\neg(P \wedge Q) \rightarrow (\neg P \vee \neg Q))$	conditional derivation, 1-13
15. $(\neg P \vee \neg Q)$	assumption for conditional derivation
16. $\neg\neg(P \wedge Q)$	assumption for indirect derivation
17. $(P \wedge Q)$	double negation 16
18. P	simplification, 17
19. $\neg\neg P$	double negation, 18
20. $\neg Q$	modus tollendo ponens, 15, 19
21. Q	simplification, 17
22. $\neg(P \wedge Q)$	indirect derivation 16-21
23. $((\neg P \vee \neg Q) \rightarrow \neg(P \wedge Q))$	conditional derivation, 15-22
24. $(\neg(P \wedge Q) \leftrightarrow (\neg P \vee \neg Q))$	bicondition,14, 23

9.7 Using theorems

Every sentence of our logic is, in semantic terms, one of three kinds. It is either a tautology, a contradictory sentence, or a contingent sentence. We have already defined "tautology" (a sentence that must be true) and "contradictory sentence" (a sentence that must be false). A contingent sentence is a sentence that is neither a tautology nor a contradictory sentence. Thus, a contingent sentence is a sentence that might be true, or might be false.

Here is an example of each kind of sentence:

(Pv¬P)

(P↔¬P)

P

The first is a tautology, the second is a contradictory sentence, and the third is contingent. We can see this with a truth table.

P	¬P	(Pv¬P)	(P↔¬P)	P
T	*F*	*T*	*F*	*T*
F	*T*	*T*	*F*	*F*

Notice that the negation of a tautology is a contradiction, the negation of a contradiction is a tautology, and the negation of a contingent sentence is a contingent sentence.

¬(Pv¬P)

¬(P↔¬P)

¬P

P	¬P	(Pv¬P)	¬(Pv¬P)	(P↔¬P)	¬(P↔¬P)
T	*F*	*T*	*F*	*F*	*T*
F	*T*	*T*	*F*	*F*	*T*

A moment's reflection will reveal that it would be quite a disaster if either a contradictory sentence or a contingent sentence were a theorem of our propositional logic. Our logic was designed to produce only valid arguments. Arguments that have no premises, we observed, should have conclusions that must be true (again, this follows because a sentence that can be proved with no premises could be proved with any premises, and so it had better be true no matter what premises we use). If a theorem were contradictory, we would know that we could prove a falsehood. If a theorem were contingent, then sometimes we could prove a falsehood (that is, we could prove a sentence that is under some conditions false). And, given that we have adopted indirect derivation as a proof method, it follows that once we have a contradiction or a contradictory sentence in an argument, we can prove anything.

Theorems can be very useful to us in arguments. Suppose we know that neither Smith nor Jones will go to London, and we want to prove, therefore, that Jones will not go to London. If we allowed ourselves to use one of De Morgan's theorems, we could make quick work of the argument. Assume the following key.

P: Smith will go to London.

Q: Jones will go to London.

And we have the following argument:

1. $\neg(P \vee Q)$	premise
2. $(\neg(P \vee Q) \leftrightarrow (\neg P \wedge \neg Q))$	theorem
3. $(\neg P \wedge \neg Q)$	equivalence, 2, 1
4. $\neg Q$	simplification, 3

This proof was made very easy by our use of the theorem at line 2.

There are two things to note about this. First, we should allow ourselves to do this, because if we know that a sentence is a theorem, then we know that we could prove that theorem in a subproof. That is, we could replace line 2 above with a long subproof that proves **(¬(P v Q)↔(¬P ^ ¬Q))**, which we could then use. But if we are certain that **(¬(P v Q)↔(¬P ^ ¬Q))** is a theorem, we should not need to do this proof again and again, each time that we want to make use of the theorem.

The second issue that we should recognize is more subtle. There are infinitely many sentences of the form of our theorem, and we should be able to use those also. For example, the following sentences would each have a proof identical to our proof of the theorem **(¬(P v Q)↔(¬P ^ ¬Q))**, except that the letters would be different:

(¬(R v S) ↔ (¬R ^ ¬S))

(¬(T v U) ↔ (¬T ^ ¬U))

(¬(V v W) ↔ (¬V ^ ¬W))

This is hopefully obvious. Take the proof of **(¬(P v Q)↔(¬P ^ ¬Q))**, and in that proof replace each instance of **P** with **R** and each instance of **Q** with **S**, and you would have a proof of **(¬(R v S)↔(¬R ^ ¬S))**.

But here is something that perhaps is less obvious. Each of the following can be thought of as similar to the theorem **(¬(P v Q)↔(¬P ^ ¬Q))**.

(¬((P^Q) v (R^S))↔(¬(P^Q) ^ ¬(R^S)))

(¬(T v (Q v V))↔(¬T ^ ¬(Q v V))

(¬((Q↔P) v (¬R→¬Q))↔(¬(Q↔P) ^ ¬(¬R→¬Q)))

For example, if one took a proof of **(¬(P v Q)↔(¬P ^ ¬Q))** and replaced each initial instance of **P** with **(Q↔P)** and each initial instance of **Q** with **(¬R→¬Q)**, then one would have a proof of the theorem **(¬((Q↔P) v (¬R→¬Q))↔(¬(Q↔P) ^ ¬(¬R→¬Q)))**.

We could capture this insight in two ways. We could state theorems of our metalanguage and allow that these have instances. Thus, we could take **(¬(Φ v Ψ) ↔ (¬Φ ^ ¬Ψ))** as a metalanguage theorem, in which we could replace each **Φ** with a sentence and each **Ψ** with a sentence and get a particular instance of a theorem. An alternative is to allow that from a theorem we can produce other theorems through substitution. For ease, we will take this second strategy.

Our rule will be this. Once we prove a theorem, we can cite it in a proof at any time. Our justification is that the claim is a theorem. We allow substitution of any atomic sentence in the theorem with any other sentence if and only if we replace each initial instance of that atomic sentence in the theorem with the same sentence.

Before we consider an example, it is beneficial to list some useful theorems. There are infinitely many theorems of our language, but these ten are often very helpful. A few we have proved. The others can be proved as an exercise.

T1 $(P \vee \neg P)$

T2 $(\neg(P \rightarrow Q) \leftrightarrow (P \wedge \neg Q))$

T3 $(\neg(P \vee Q) \leftrightarrow (\neg P \wedge \neg Q))$

T4 $((\neg P \vee \neg Q) \leftrightarrow \neg(P \wedge Q))$

T5 $(\neg(P \leftrightarrow Q) \leftrightarrow (P \leftrightarrow \neg Q))$

T6 $(\neg P \rightarrow (P \rightarrow Q))$

T7 $(P \rightarrow (Q \rightarrow P))$

T8 $((P \rightarrow (Q \rightarrow R)) \rightarrow ((P \rightarrow Q) \rightarrow (P \rightarrow R)))$

T9 $((\neg P \rightarrow \neg Q) \rightarrow ((\neg P \rightarrow Q) \rightarrow P))$

T10 $((P \rightarrow Q) \rightarrow (\neg Q \rightarrow \neg P))$

Some examples will make the advantage of using theorems clear. Consider a different argument, building on the one above. We know that neither is it the case that if Smith goes to London, he will go to Berlin, nor is it the case that if Jones goes to London he will go to Berlin. We want to prove that it is not the case that Jones will go to Berlin. We add the following to our key:

R: Smith will go to Berlin.

S: Jones will go to Berlin.

And we have the following argument:

1.	$\neg((P \to R) \lor (Q \to S))$	premise
2.	$(\neg((P \to R) \lor (Q \to S)) \leftrightarrow (\neg(P \to R) \land \neg(Q \to S)))$	theorem T3
3.	$(\neg(P \to R) \land \neg(Q \to S))$	equivalence, 2, 1
4.	$\neg(Q \to S)$	simplification, 3
5.	$(\neg(Q \to S) \leftrightarrow (Q \land \neg S))$	theorem T2
6.	$(Q \land \neg S)$	equivalence, 5, 4
7.	$\neg S$	simplification, 6

Using theorems made this proof much shorter than it might otherwise be. Also, theorems often make a proof easier to follow, since we recognize the theorems as tautologies—as sentences that must be true.

9.8 Problems

1. Prove each of the following arguments is valid.

    a. Premises: **((P^Q) ↔ R), (P ↔ S), (S ^ Q)**. Conclusion: **R**.
    b. Premises: **(P ↔ Q)**. Conclusion: **((P → Q) ^ (Q → P))**.
    c. Premises: **P, ¬Q**. Conclusion: **¬(P ↔ Q)**.
    d. Premises: **(¬PvQ), (Pv¬Q)**. Conclusion: **(P ↔ Q)**.
    e. Premises: **(P ↔ Q), (R ↔ S)**. Conclusion: **((P^R) ↔ (Q^S))**.
    f. Premises: **((PvQ) ↔ R), ¬(P ↔ Q)**. Conclusion: **R**.
    g. Conclusion: **((P ↔ Q) ↔ (¬P ↔ ¬Q))**.
    h. Conclusion: **((P → Q) ↔ (¬P v Q))**.

2. Prove each of the following theorems.

    a. T2
    b. T3
    c. T5
    d. T6
    e. T7
    f. T8
    g. T9

h. $((P \wedge Q) \leftrightarrow \neg(\neg P \vee \neg Q))$
i. $((P \rightarrow Q) \leftrightarrow \neg(P \wedge \neg Q))$

3. Here are some passages from literature, philosophical works, and important political texts. Hopefully you recognize some of them. Find the best translation into propositional logic. Because these are from diverse texts you will find it easiest to make a new key for each sentence.

   a. "Neither a borrower nor a lender be." (Shakespeare, *Hamlet*.)
   b. "My copy-book was the board fence, brick wall, and pavement." (Frederick Douglass, *Narrative of the Life of Frederick Douglass*.)
   c. "The bourgeoisie has torn away from the family its sentimental veil, and has reduced the family relation to a mere money relation." (Marx and Engels, *The Communist Manifesto*.)
   d. "The Senate shall chuse their other Officers, and also a President pro tempore, in the Absence of the Vice President, or when he shall exercise the Office of President of the United States." (The Constitution of the United States.)
   e. "Excessive bail shall not be required, nor excessive fines imposed, nor cruel and unusual punishments inflicted." (The Constitution of the United States.)
   f. "Annual income twenty pounds, annual expenditure nineteen nineteen and six, result happiness. Annual income twenty pounds, annual expenditure twenty pounds ought and six, result misery." (Charles Dickens, *Great Expectations*.)
   g. "Thou shalt get kings, though thou be none." (Shakespeare, *Macbeth*.)
   h. "If a faction consists of less than a majority, relief is supplied by the republican principle, which enables the majority to defeat its sinister views by regular vote." (Federalist Papers.)

4. In normal colloquial English, write your own valid argument with at least two premises, at least one of which is a biconditional. Your argument should just be a paragraph (not an ordered list of sentences or anything else that looks like formal logic). Translate it into propositional logic and prove it is valid.

5. In normal colloquial English, write your own valid argument with at least two premises, and with a conclusion that is a biconditional. Your argument

should just be a paragraph (not an ordered list of sentences or anything else that looks formal like logic). Translate it into propositional logic and prove it is valid.

6. Here is a passage from Aquinas's reflections on the law, The Treatise on the Laws. Symbolize this argument and prove it is valid.

> A law, properly speaking, regards first and foremost the order to the common good. Now if a law regards the order to the common good, then its making belongs either to the whole people, or to someone who is the viceregent of the whole people. And therefore the making of a law belongs either to the whole people or to the viceregent of the whole people.

---

[11] From Hume's *Enquiry Concerning Human Understanding,* p.161 in Selby-Bigge and Nidditch (1995 [1777]).

10. Summary of Propositional Logic

10.1 Elements of the language

- Principle of Bivalence: each sentence is either true or false, never both, never neither.
- Each atomic sentence is a sentence.
- Syntax: if **Φ** and **Ψ** are sentences, then the following are also sentences
    - **¬Φ**
    - **(Φ→Ψ)**
    - **(Φ ^ Ψ)**
    - **(Φ v Ψ)**
    - **(Φ↔Ψ)**
- Semantics: if **Φ** and **Ψ** are sentences, then the meanings of the connectives are fully given by their truth tables. These truth tables are:

Φ	¬Φ
T	*F*
F	*T*

Φ	Ψ	(Φ→Ψ)
T	*T*	*T*
T	*F*	*F*
F	*T*	*T*
F	*F*	*T*

Φ	Ψ	(Φ ^ Ψ)
T	*T*	*T*
T	*F*	*F*
F	*T*	*F*
F	*F*	*F*

Φ	Ψ	(Φ v Ψ)
T	*T*	*T*
T	*F*	*T*
F	*T*	*T*
F	*F*	*F*

Φ	Ψ	(Φ↔Ψ)
T	*T*	*T*
T	*F*	*F*
F	*T*	*F*
F	*F*	*T*

- A sentence of the propositional logic that must be true is a tautology.
- A sentence that must be false is a contradictory sentence.
- A sentence that is neither a tautology nor a contradictory sentence is a contingent sentence.
- Two sentences **Φ** and **Ψ** are equivalent, or logically equivalent, when **(Φ↔Ψ)** is a theorem.

10.2 Reasoning with the language

- An argument is an ordered list of sentences, one sentence of which we call the "conclusion" and the others of which we call the "premises".
- A valid argument is an argument in which: necessarily, if the premises are true, then the conclusion is true.
- A sound argument is a valid argument with true premises.
- Inference rules allow us to write down a sentence that must be true, assuming that certain other sentences are true. We say that the new sentence is "derived from" those other sentences using the inference rule.
- Schematically, we can write out the inference rules in the following way (think of these as saying, if you have written the sentence(s) above the line, then you can write the sentence below the line):

Modus ponens	Modus tollens	Double negation	Double negation
$(\Phi \rightarrow \Psi)$ Φ _____ Ψ	$(\Phi \rightarrow \Psi)$ $\neg\Psi$ _____ $\neg\Phi$	Φ _____ $\neg\neg\Phi$	$\neg\neg\Phi$ _____ Φ
Addition	**Addition**	**Modus tollendo ponens**	**Modus tollendo ponens**
Φ _____ $(\Phi \vee \Psi)$	Ψ _____ $(\Phi \vee \Psi)$	$(\Phi \vee \Psi)$ $\neg\Phi$ _____ Ψ	$(\Phi \vee \Psi)$ $\neg\Psi$ _____ Φ
Adjunction	**Simplification**	**Simplification**	**Bicondition**
Φ Ψ _____ $(\Phi \wedge \Psi)$	$(\Phi \wedge \Psi)$ _____ Φ	$(\Phi \wedge \Psi)$ _____ Ψ	$(\Phi \rightarrow \Psi)$ $(\Psi \rightarrow \Phi)$ _____ $(\Phi \leftrightarrow \Psi)$
Equivalence	**Equivalence**	**Equivalence**	**Equivalence**
$(\Phi \leftrightarrow \Psi)$ Φ _____ Ψ	$(\Phi \leftrightarrow \Psi)$ Ψ _____ Φ	$(\Phi \leftrightarrow \Psi)$ $\neg\Phi$ _____ $\neg\Psi$	$(\Phi \leftrightarrow \Psi)$ $\neg\Psi$ _____ $\neg\Phi$

- A proof (or derivation) is a syntactic method for showing an argument is valid. Our system has three kinds of proof (or derivation): direct, conditional, and indirect.
- A direct proof (or direct derivation) is an ordered list of sentences in which every sentence is either a premise or is derived from earlier lines using an inference rule. The last line of the proof is the conclusion.
- A conditional proof (or conditional derivation) is an ordered list of sentences in which every sentence is either a premise, is the special assumption for conditional derivation, or is derived from earlier lines using an inference rule. If the assumption for conditional derivation is Φ, and we derive as some step in the proof Ψ, then we can write after this $(\Phi \rightarrow \Psi)$ as our conclusion.

- An indirect proof (or indirect derivation, and also known as a reductio ad absurdum) is: an ordered list of sentences in which every sentence is either 1) a premise, 2) the special assumption for indirect derivation (also sometimes called the “assumption for reductio”), or 3) derived from earlier lines using an inference rule. If our assumption for indirect derivation is **¬Φ**, and we derive as some step in the proof **Ψ** and also as some step of our proof **¬Ψ**, then we conclude that **Φ**.
- We can use Fitch bars to write out the three proof schemas in the following way:

$$\begin{array}{|l} \phi \\ \vdots \\ \psi \end{array}$$

$$\begin{array}{r|l} & \begin{array}{|l} \phi \\ \hline \vdots \\ \psi \end{array} \\ \rhd & (\phi \to \psi) \end{array}$$

$$\begin{array}{r|l} & \begin{array}{|l} \neg\phi \\ \hline \vdots \\ \psi \\ \neg\psi \end{array} \\ \rhd & \phi \end{array}$$

- A sentence that we can prove without premises is a theorem.
- Suppose **Φ** is a theorem, and it contains the atomic sentences $P_1 \ldots P_n$. If we replace each and every occurrence of one of those atomic sentences P_i in **Φ** with another sentence **Ψ**, the resulting sentence is also a theorem. This can be repeated for any atomic sentences in the theorem.