# Chapter 2

# Logic

## 2.1 The Rain vs. The Water Truck

If you wake up in the morning and observe that the road outside is wet, you are likely to conclude that it rained during the night. That you did not see it rain is neither here nor there. You know that rain is something that happens from time to time, and that rain has the effect of causing roads to become wet. No doubt you have seen it rain in the past and recall that it causes the roads to appear substantially as they do now.

But now suppose your roommate enters the room and claims that actually the wetness of the road was caused by a recently cleared automobile accident involving a water truck. Confronted with these differing explanations, how are you to determine which is true?

Let us suppose for a moment that the water truck theory is correct. You might reason that such a momentous accident would surely be a topic of discussion on the local radio station, and turn on the radio accordingly. It turns out that there is no mention of such an accident. The water truck theory is thus seen to false, and the rain theory seems more likely.

Not so fast. The reasoning described in the previous paragraph involved three statements. The first was an if-then statement: If there was a car accident involving a water truck on my street, then there will be some mention of it on the local radio station. The second was a statement of fact: There was no mention of such an accident on the radio. The third was a conclusion: Therefore, the wetness of the road was not caused by a recent water truck accident. If the first two statements are true, then the third must be true as well. Consequently, your reasoning is logically valid.

But are the first two statements true? You might argue that the second statement, at least, is easily verified. You and your roommate can listen to the radio for a time and come to some agreement as to whether any mention of the accident is made. The first statement is different. Your roommate might claim that the accident happened too recently to be mentioned on the local news, or that the announcers misplaced the item relating to the accident and thereby neglected to read it on the air. If your roommate is correct then your argument is robbed of its force; while still logically valid, it is not logically sound.

You might reply that since there is evidently no demolished water truck on the road right now, it follows that the accident happened at least thirty minutes ago. It would take at least that long for such an accident to be cleared, you see. Since thirty minutes is sufficient time for the radio station to have heard about it, your roommate's first reply is not correct. In so replying, of course, you have made another logical argument. You might then go on to add that the radio announcers were not likely to misplace such an important news item, and consequently the most reasonable explanation for the non-mention of the accident is the non-happening of the accident in the first place.

To which your roommate might reply that the local radio station is staffed by inefficient morons who take forever to report important stories. Since your best friend works for the radio station you take offense at this remark, and the whole thing dissolves into a shouting match.

Still, it would be nice to know what really caused the wetness of the road.

Let us now suppose that you successfully persuade your roommate that your if-then statement was true and that the water truck theory is unlikely to be correct. Is this evidence for the correctness of the rain theory? As a practical matter it probably is; there are only two explanations under consideration and plausible third alternatives are not forthcoming. If one of the two theories is shown to be false then it is not unreasonable to conclude that the other is true.

Logically the situation is quite different. There is an almost endless supply of theories to explain the wetness of the road: Perhaps a group of pranksters dumped water on the road for mischievous reasons of their own. Maybe a large pack of ducks flew by recently and their combined drippings created the illusion of rain. Possibly ghosts or poltergeists are responsible for the water on the road. These explanations might seem implausible, but they are not impossible and, logically speaking, that is all that matters. So unless you are prepared to say that the collapse of the water truck theory constitutes evidence in favor of the poltergeist theory, you can not translate the failure of your roommate's idea into the success of your own.

What you *can* do is think about further consequences of the rain theory and then check to see if those consequences hold. Ideally you would look for consequences that are unique to the rain theory. For example, if it rained you might expect the covering of water on the road to be relatively uniform. You would also expect that the leaves of nearby trees would be covered with water. If there are cars parked on the street then you would expect that the roofs of the cars would be wet but the pavement beneath the cars would be dry. The water truck theory, by contrast, would lead us to expect a large amount of water in the vicinity of the accident, but very little water a short distance away. A bursting water truck would have caused water to run underneath any parked cars, and this pavement would be wet as well.

All of these predictions could have been expressed as if-then statements, just as we did previously. Now suppose that in each case the prediction made by the rain theory is born out. We would then say that we have amassed a considerable amount of evidence in favor of the rain theory and against those rival theories having different predictive consequences. But is there any point at which we are justified in saying that we have proved to a logical certainty that the rain theory is correct?

Alas, there is not. To see this let us suppose that we have a statement of the following form: If the rain theory is correct, then we should observe $X$ to be true. If we subsequently observe that $X$ is indeed true, we are led to the following argument:

- If the rain theory is correct, then we should observe $X$ to be true.
- $X$ is observed to be true. Therefore,
- The rain theory is correct.

Claiming to have proven the rain theory to be correct is equivalent to saying that the argument above is valid. That would mean that if the first two statements are true, then the third statement must be true as well. The trouble is that though $X$ may be a logical consequence of the rain theory, it is surely a logical consequence of numerous other theories as well. Thus the fact that $X$ is true and $X$ is a logical consequence of the rain theory does not guarantee that the rain theory is correct. The argument above is not logically valid.

It follows that while we may amass sufficient evidence to convince any reasonable person that it rained last night and that we might conclusively disprove the rain theory by showing that one of its logical consequences fails to hold, we never reach a point where the rain theory can be said to be logically certain.

A dissatisfying conclusion, perhaps, but no less true for that.

## 2.2 Deductive and Inductive Reasoning

Logic is what helps us distinguish good arguments from bad arguments. Since proper argumentation is especially important in science we will briefly discuss the nature of scientific reasoning.

It is commonly believed that science is a grand search for the truth about nature. The trouble is that truth is a slippery beast. No matter how much data we assemble in support of a theory, it is always logically possible that we are mistaken. Nonetheless, scientists do not hesitate to describe their theories as true. It follows they have something other than logical certainty in mind when they do so.

Scientists are routinely confronted with competing theories for explaining experimental data. To distinguish between them, they avail themselves of two sorts of reasoning. In the first we observe that Theory X has successfully predicted the outcomes of numerous experiments, and conclude that it will continue to work for experiments not yet performed. This is known as inductive reasoning, and it is fraught with peril. Extrapolating from particular instances to general principles can lead to false conclusions. To use a famous, if frivolous example: Suppose you are a turkey. You notice that for 364 consecutive days you are fed at 7:00 in the morning, and from this evidential base you conclude that on the 365th day you will once again be fed at that time. But you would be wrong. For the 365th day turns out to be Thanksgiving.

Our hypothetical turkey notwithstanding, it sure feels rational to begin with the premise that the Sun has risen each morning for every day of recorded history and conclude that it will rise again tomorrow. The problem of explaining the effectiveness of inductive reasoning has plagued philosophers for centuries, but we will be content to observe that it is, in fact, effective.

In the second sort of reasoning we make certain assumptions about the world, and argue that if these assumptions are true then certain other things must be true as well. This is known as deductive reasoning, and certainly seems a bit more solid. No one, upon hearing that "All men are mortal," and "Socrates was a man," will hesitate to conclude that "Socrates was mortal." As a practical matter, however, presenting a deductive argument merely forces you to defend the assumptions that went into it, and that is often difficult to do.

The word "inductive" derives from Latin words meaning "to lead into". The idea is that when reasoning inductively we use a handful of specific instances to lead into a more general conclusion. By contrast, "deductive" comes from Latin words meaning "to lead down from". Here, the idea is that deductive reasoning involves moving down from given assumptions to their necessary consequences. In mathematics we are exclusively interested in deductive reasoning. Before discussing that, however, let us consider an illustrative example of both deductive and inductive reasoning in action.

Toward the end of the seventeenth century Isaac Newton devised his laws of mechanics. By judicious application of these laws it was possible to determine the trajectories of various moving objects merely by knowing certain easily measured quantities related to them. For example, if I throw a tennis ball straight upward with some known initial velocity, Newton's laws allow me to determine the highest point the ball will reach, the time it will take for it to reach that point, the time it will take for the ball to hit the ground, and the speed it will be traveling when it does so. In other words, deductive reasoning allowed scientists to argue that if Newton's theories were correct, then moving objects should follow particular paths under experimental conditions. When numerous experiments showed that actual moving objects routinely behaved precisely as Newton's theories said they should, these theories won the acceptance of the scientific community.

Since Newton's theories worked so well for moving objects in the laboratory, scientists reasoned inductively to conclude that they would also work for larger objects like planets and moons. Astronomers, now reasoning deductively, used Newton's laws to predict the trajectories of the planets around the Sun, and the excellent fit of data to theory was taken as further vindication for Newton.

But then it was discovered that the orbit of Uranus, the most distant planet known at the time, deviated measurably from the path predicted by Newton's laws. This bit of data had to be explained, and several possibilities came to mind.

Perhaps Newton was simply wrong. Since his theories had already proven their worth in countless experiments this seemed unlikely.

More reasonable was that Newton's theories were correct when applied to nearby objects, but broke down when applied to objects separated by vast, astronomical distances. In other words, the inductive extrapolation from particular instances (the successful predictions of laboratory experiments) to general principles (Newton's laws are also correct for planets) might not have been valid. While this would explain why the observed path of Uranus differed from the experimentally predicted one, it is not very satisfying. Simply declaring Newton's laws to be false over vast distances helps neither in determining the correct path of Uranus, nor in determining the motions of any other astronomical bodies we might subsequently discover.

Happily, there was a third possibility. Perhaps the planets were moving just as Newton's laws predicted, but there was an unknown planet beyond Uranus that had not yet been discovered. A proper determination of the orbit of Uranus would then require a consideration of the motion of this hypothetical planet.

Assume for the moment such a planet exists. How would we find it? Randomly searching the sky with a telescope was unpromising. After all, people had been using telescopes to investigate the heavens for quite some time, and this hypothetical planet had so far eluded discovery. A better approach would be to use Newton's laws to determine, deductively, what properties our planet must have to account for Uranus' bizarre motion. This would give us a strong indication of where in the sky to look.

And that is precisely what happened. In 1846, astronomer John Herschel, using theoretical calculations based upon Newton's laws, successfully observed the planet Neptune. It was a fine day in the history of astronomy and one more vindication for Newton's laws.

Nonetheless, it was subsequently shown by Albert Einstein that Newton's theories were wrong in fundamental ways. We will discuss this further in the next chapter.

Thus, when scientists say a theory is true, they do not mean that it is logically certain. Rather, they mean that it has passed so many experimental tests that it would be mere stubbornness to deny it. Of course, successfully predicting the outcomes of one thousand experiments tells you nothing about the outcome of the thousand and first. For this reason no scientific theory is ever one hundred percent certain. Instead it is established to some high degree of certainty.

I should point out that in using the term "experiment" I am not necessarily referring to a contrived situation in a laboratory setting. I mean simply that a scientific theory must have predictive consequences that can be checked against nature's data. Theories from which nothing testable can be deduced are unscientific, which is different from saying that they are false. For example, suppose I believe the wetness of the road is the result of poltergeists. What deductions can I make from this starting point? Sadly, there are none. There is no obviously true statement that begins with "If the poltergeist theory is correct," and ends with "then I should observe $X$." The poltergeist theory is unscientific.

Of course, very big books get written on the philosophical topics discussed above. What I have described is the way science ought to work, and in its finest moments does work. But let us not have the impression that scientists are perfect reasoning machines, deciding their disputes solely by an objective consideration of nature's data. The history and philosophy of science are fascinating subjects in their own right, but they are not the subject of this book.

## 2.3 Vocabulary

Several bits of logical terminology have arisen naturally in our previous discussions. Allow me to call your attention to them.

First, by a ***proposition*** we mean any statement that can reasonably be said to be either true or false, but not both. "It rained last night," and "The water on the road is the work of poltergeists," are examples of propositions. Statements like "Listen to me!" and "Are you going to the party?" are not propositions.

Distinguishing propositions from non-propositions has nothing to do with our ability to determine the truth or falsity of anything. The statement, "There exists life on other planets" is a proposition. It is either true or false, our current inability to decide which notwithstanding.

It is customary to denote arbitrary propositions by capital letters. Thus, when you see something like "Let $X$ be a proposition...", you should interpret that to mean that $X$ is a statement that makes some factual assertion.

For some statements it is ambiguous whether they are propositions or not. It seems clear that statements of opinion are not propositions. The statement, "Ice cream tastes good" is arguably a matter of opinion, and therefore can not be said to be either true or false. What about moral claims? If I say "Murder is wrong," am I merely stating an opinion or am I making an assertion of fact? If you think the answer is obvious, ask the same question of a more controversial moral assertion.

And what about category errors? This is a philosophical term referring to a statement that ascribes a property to something that can not possibly have that property. If I say "My opinions are green" or "The number five went to the movies" I am making a category error. Opinions can not have colors and numbers can not go to movies, you see. Are such statements false, and therefore propositions (opinions are, in fact, not green, for they have no color at all)? Or are they meaningless; an abuse of language that makes no meaningful assertion at all (in which case they certainly can not be said to be either true or false)?

Happily, these are philosophical thickets that need not detain us. In the context of mathematics it is perfectly clear what is a proposition and what is not.

An ***argument*** is a series of propositions (called ***premises***) followed by one additional proposition (called the ***conclusion***). If the conclusion must be true whenever all of the premises themselves are true, then the argument is said to be ***valid***. If in addition the premises really are true, then the argument is said to be ***sound***.

Validity and soundness are two different notions, and it is important that you understand the difference between them. Every sound argument is valid, but not every valid argument is sound.

For example, in the argument:

- All healthy dogs have four legs.
- Spot has four legs. Therefore,
- Spot is a dog,

the premises are the statements "All dogs have four legs" and "Spot has four legs". Sadly, this argument is not valid. Even if both of these premises are true, the conclusion that "Spot is a dog" is not warranted. Spot might be a giraffe, or a badly injured spider, or a table. More precisely, we might say that the proposition "All healthy dogs have four legs" is not equivalent to the proposition "All four-legged things are healthy dogs."

The argument:

- All prime numbers are even.
- 5 is a prime number. Therefore,
- 5 is even,

is valid. If the premises are true, then the conclusion must be true as well. But since the first premise is false, the argument is not sound.

Finally the argument:

- All prime numbers other than 2 are odd.
- 5 is a prime number. Therefore,
- 5 is odd.

is sound. Both of the premises are true, and the conclusion follows from them.

By the way, do not get the impression that every argument consists of two premises followed by a conclusion. We will consider more complex arguments as we go along. Arguments like the ones above, consisting of two premises (sometimes further divided into a ***major premise*** and a ***minor premise***) and a conclusion are sometimes referred to as ***syllogisms***.

Given any proposition $X$ we can identify a second proposition that asserts that $X$ is false. This second statement is referred to as the ***negation*** of $X$, or more simply as "not $X$". Thus, if $X$ is the proposition "It snowed yesterday," then not $X$ is the statement "It did not snow yesterday." If $X$ is the proposition "Lucy is five years old," then not $X$ is the statement "Lucy is not five years old." If $X$ is the proposition "All crows are black," then not $X$ is the statement "There is at least one crow that is not black" (think about that one for a while).

We conclude this section by mentioning two important principles. A proposition $X$ can not be simultaneously true and false. This is known as ***The Law of Non-Contradiction***. Also, given any proposition $X$, either $X$ or not $X$ is true. This is known as ***The Law of the Excluded Middle***.

## 2.4 The Rules

Now let us get down to the business of establishing the basic rules of formal logic.

The first thing to realize is that logic is interested in validity, not soundness. Just as you do not use a hammer to slice an apple pie, you do not use logic to determine the truth or falsity of anything. A person who claims the moon is made of green cheese is not behaving illogically. He is merely mistaken on an issue of fact.

We have already introduced the negation of a proposition $X$. Since it is cumbersome to write "The negation of $X$" or "not $X$", we will use the abbreviation $\neg X$.

Actually, this may be a good time to mention that mathematical symbols are simply abbreviations for English words. That is all. They are not hieroglyphics. When you are reading a piece of technical mathematics, develop the habit of mentally translating the symbols into common English words.

### 2.4.1 Not

If $X$ is a proposition and $X$ is true, then $\neg X$ is false. And if $X$ is false, then $\neg X$ is true. We can capture that information in the following nifty graphic, known portentously as a ***Truth Table***:

| $X$ | $\neg X$ |
|---|---|
| T | F |
| F | T |

This table is read as follows: The first row below the horizontal line tells us that if the proposition $X$ is true then the proposition $\neg X$ is false. The next line tells us that if $X$ is false then $\neg X$ is true. Since this covers all possibilities we can move on.

### 2.4.2 If-Then

Some propositions are formed by combining two or more other propositions. The if-then statements we considered previously are examples of this. If $X$ and $Y$ are propositions then the statement "If $X$ then $Y$" is called a ***conditional***. It is abbreviated $X \to Y$. The proposition $X$ is called the

***antecedent*** of the conditional, and the proposition $Y$ is called the ***conclusion***.

Conditional statements are especially important in mathematics, as we shall see in the next chapter. You can think of the statement $X \rightarrow Y$ as saying that $X$ implies $Y$, or that $Y$ follows as a logical consequence of $X$.

We now must decide when this statement is true and when it is false. The truth of the compound statement "If $X$ then $Y$" will certainly depend on the truth or falsity of $X$ and $Y$ taken individually. To make things concrete, let us suppose that $X$ is the proposition "You build it" and $Y$ is the proposition "They will come". Then we wish to determine the circumstances under which the statement "If you build it then they will come" is true.

If I build it and they do, in fact, come, then the statement is true. In other words, if both $X$ and $Y$ are true, then the conditional statement $X \rightarrow Y$ is also true. Similarly, if I build it and they do not come, then the statement $X \rightarrow Y$ is false.

So far everything is in accord with common sense. But complications arise if we suppose $X$ to be false. The trouble is that while our conditional tells us what to expect when $X$ is true, it says nothing about what happens when $X$ is false. If I don't build it, then they may or may not come for all the information my conditional statement provides.

In light of this, it is tempting to leave the truth value of $X \rightarrow Y$ undetermined in this case. Sadly, this runs afoul of the law of the excluded middle which requires us to assign a truth value to this statement. Thus, we will adopt the following convention: If you don't build it, then whether they come or not is immaterial. I have not, strictly speaking, lied to you, and therefore we will declare the conditional to be true. In other words, if $X$ is false then the conditional $X \rightarrow Y$ is true regardless of whether or not $Y$ is true.

We encapsulate this information in the following truth table:

| $X$ | $Y$ | $X \rightarrow Y$ |
|---|---|---|
| T | T | T |
| T | F | F |
| F | T | T |
| F | F | T |

Our convention has the amusing consequence that a statement such as "If Santa Claus exists then I can leap tall buildings in a single bound" are true. This is something you must think about until it ceases to seem weird.

When solving a complex mathematical problem it is often useful to proceed by breaking it up into several, simpler problems. A nice illustration of this principle is provided by the problem of constructing the truth table for a complex proposition.

As an example, consider the proposition $X \to (\neg Y)$, which is read "If $X$ then not $Y$". The truth or falsity of this proposition depends on the truth or falsity of the simpler propositions $X$ and $\neg Y$. The proposition $\neg Y$ is itself a proposition whose truth value depends on $Y$. Thus, in constructing the truth table for this proposition we will include one column for each of the individual propositions out of which our complex statement is built; namely $X$, $Y$, $\neg Y$ and $X \to (\neg Y)$. The resulting table looks like this:

| $X$ | $Y$ | $\neg Y$ | $X \to (\neg Y)$ |
|---|---|---|---|
| T | T | F | F |
| T | F | T | T |
| F | T | F | T |
| F | F | T | T |

Notice that in filling in the fourth column we were able to ignore completely the second column of the table. After all, the fourth column is a conditional statement whose antecedent is given in the first column and whose conclusion is given in the third.

### 2.4.3 And

Sometimes we wish to combine two propositions by the word "and", as in "Ducks have webbed feet and cool people study math." Such a statement is called a ***conjunction***. If $X$ and $Y$ are propositions then the conjunction "$X$ and $Y$" is abbreviated $X \wedge Y$. This symbol is easier to remember if you note its resemblance to a capital letter "A".

"And" is a sufficiently common English word that no difficulties arise in constructing its truth table. The statement $X \wedge Y$ is true only when both $X$ and $Y$ are themselves true. Otherwise the conjunction is false. Consequently, the truth table looks like this:

| $X$ | $Y$ | $X \wedge Y$ |
|---|---|---|
| T | T | T |
| T | F | F |
| F | T | F |
| F | F | F |

### 2.4.4 Or

When two propositions are combined with the word "or" the resulting proposition is referred to as a ***disjunction***. If $X$ and $Y$ are propositions, then the proposition "$X$ or $Y$" is abbreviated "$X \vee Y$".

Filling in the truth table for such a proposition may seem a simple task, but there is a difficulty lurking beneath the surface. In everyday English the word "or" is ambiguous. If I say "I will purchase a Honda or I will purchase a Toyota" the implication is that I am not going to buy both. On the other hand, if I say "In order to take an advanced math class you must have had multivariable calculus or linear algebra," I certainly do not mean to exclude those who have taken both.

In other words, in everyday speech the proposition "$X \vee Y$" can mean either "$X$ or $Y$ or both" or "$X$ or $Y$ but not both." Therefore, we must make a decision about which of these we intend when we write $X \vee Y$. Among mathematicians it is customary to use the first interpretation, and we will follow that custom. Consequently, the disjunction $X \vee Y$ will be false only when $X$ and $Y$ are both false. Its truth table will look like this:

| $X$ | $Y$ | $X \vee Y$ |
|---|---|---|
| T | T | T |
| T | F | T |
| F | T | T |
| F | F | F |

Now that we have enlarged our logical vocabulary, let us try constructing a truth table for the proposition:

$$(X \wedge Y) \to \neg((\neg X) \vee (\neg Y)).$$

To simplify things I will refer to this long proposition simply as $Z$. Determining the truth or falsity of $Z$ by plugging in random truth values for $X$ and $Y$ is a difficult task. The key is to break $Z$ up into its component pieces.

Ultimately, $Z$ is a conditional whose antecedent is $(X \wedge Y)$ and whose conclusion is

$$\neg((\neg X) \vee (\neg Y)).$$

The conclusion is itself the negation of the disjunction

$$(\neg X) \vee (\neg Y).$$

Thus, our truth table should have columns for $X$, $Y$, $X \wedge Y$, $\neg X$, $\neg Y$, $(\neg X) \vee (\neg Y)$, $\neg((\neg X) \vee (\neg Y))$ and finally $Z$. We obtain the following table:

| $X$ | $Y$ | $X \wedge Y$ | $\neg X$ | $\neg Y$ | $(\neg X) \vee (\neg Y)$ | $\neg((\neg X) \vee (\neg Y))$ | $Z$ |
|---|---|---|---|---|---|---|---|
| T | T | T | F | F | F | T | T |
| T | F | F | F | T | T | F | T |
| F | T | F | T | F | T | F | T |
| F | F | F | T | T | T | F | T |

Once again, observe that the final column of the table was filled in only by considering the third and seventh columns.

As a result of our exertions we have made a curious discovery. Regardless of the truth values of the propositions $X$ and $Y$, the proposition $Z$ is always true. Propositions such as $Z$ are referred to as ***tautologies***. Actually, this result is not so surprising if we translate $Z$ into everyday language. It is effectively saying, "If it is the case that both $X$ and $Y$ are true, then neither $X$ nor $Y$ is false." Put this way it is easy to see that $Z$ is always true.

### 2.4.5 If And Only If

We will close this section with one further method of connecting two propositions. This is known as the ***biconditional*** and is abbreviated $X \leftrightarrow Y$. This is read "$X$ if and only if $Y$". You can view this as a shorthand way of writing two distinct conditional statements. One of them is "$X \to Y$" and the other is "$Y \to X$". In fact we can use the notation we have defined in the previous sections to say that the biconditional

$$X \leftrightarrow Y$$

is equivalent to the conjunction

$$(X \to Y) \wedge (Y \to X).$$

From this we conclude that the truth table of the biconditional is:

| $X$ | $Y$ | $X \leftrightarrow Y$ |
|---|---|---|
| T | T | T |
| T | F | F |
| F | T | F |
| F | F | T |

Notice that when $X$ and $Y$ have the truth values given by the second row the conditional $X \rightarrow Y$ is false. Similarly, in the third row we have that $Y \leftarrow X$ is false.

A good way to think about biconditionals is to interpret the proposition $P \leftrightarrow Q$ as saying that $P$ and $Q$ are ***logically equivalent***. By this we mean that as soon as you know the truth value of either $P$ or $Q$, you automatically know the truth value of the other one. More generally we say that two propositions $P$ and $Q$ are logically equivalent if the biconditional $P \leftrightarrow Q$ is a tautology.

Finally, we must once again come to terms with the fact that, according to the fourth row of the table, the statement "Santa Claus exists if and only if there is a purple elephant in Africa" is true. Isn't logic fun?

### 2.4.6 Why These Rules?

The foregoing rules are easily summarized. An "and" statement is true only when both of its parts are true. An "or" statement is true unless both of its parts are false. An "if-then" is false if the first part is true and the second part is false, but it is true otherwise. An "if and only if" statement is true when both of its parts have the same truth value, and is false otherwise.

But why, you may wonder, are these the rules? Why not some other, more convenient, rules? Well, keep in mind that in defining the rules for sound reasoning we are saying nothing about how the world actually is. It is not as if we performed some experiment whose results told us the proper way of defining the rules. For example, when we say an "and" statement is true only when each of its parts is true, we are not revealing some law of nature. We are only revealing what we mean when we use the word "and". The rules given above should be construed as definitions of words, not fundamental principles. Indeed, the word "logic" comes from the Greek word *logos*, meaning "speech" or "discourse".

### 2.4.7 Other Connectives

Which raises the following question: Given any two propositions there are four ways of assigning truth values to them (either both are true, both false, or one is true and the other false). Having made those assignments, there are then sixteen different ways of assigning truth values to the resulting, conjoined proposition (in other words, sixteen different ways of filling in the third column of the truth table). Why do we have words for describing some of those sixteen ways and not for others?

For example, let us define a new connective, which we will denote "$\oplus$". This connective is defined as follows:

| $X$ | $Y$ | $X \oplus Y$ |
|---|---|---|
| T | T | F |
| T | F | T |
| F | T | T |
| F | F | F |

Actually this connective behaves exactly like an "either-or" statement; specifically, it is true only when its parts have different truth values. If we desired, we could add this to our already lengthy list of logical connectives.

Here is another one, this time denoted $\sim$:

| $X$ | $Y$ | $X \sim Y$ |
|---|---|---|
| T | T | F |
| T | F | T |
| F | T | F |
| F | F | T |

This particular connective does not correspond to any familiar English word, but it is no less legitimate for that. We could invent a word that would describe this situation, if we so desired. We do not do that only because mathematicians have not found it especially helpful to do so.

In fact, any of the sixteen possible truth tables could be used to describe some logical connective. We have isolated the ones we have only because they have turned out to be the most useful in practice.

## 2.5 Quantifiers

Now consider a proposition such as "All crows are black". To establish that this were true I would have to examine every crow in the world and verify that all of them were indeed black. In principle, I suppose, we could gather up all of the non-black things in the world and show that not one of those things is a crow, but that seems rather impractical. Of course, gathering up all the world's crows seems scarcely more promising. In practice we would examine some large number of crows, verify that all of them are black, and then reason inductively to the conclusion that, actually, all crows are black. If we were feeling cautious we might say "Every crow that has ever been observed has been black."

What if we wanted to disprove the statement? It would not be necessary to show that, actually, all crows are white, or that there are no black crows on the entire planet. That would be overkill. Actually, the statement "All crows are black" is seen to be false as soon as a single non-black crow turns up. This is a consequence of the definition of the word "all". We might even say that the negation of the statement "All crows are black" is "There exists at least one crow that is not black."

Let us make this more precise. Suppose I have some collection of objects and I make the assertion that every member of the collection has a certain property. That assertion is shown to be false by producing a single object in the collection lacking the property. For example, if I say all of the objects in my suitcase weigh less than five ounces, then a single object weighing more than five ounces would be enough to disprove my statement.

Now let's turn it around. Suppose I say "There is at least one student at Big State University who has purple hair." A single Big State student possessing purple hair would be enough to prove this statement to be true. To disprove this statement, I would have to show that every single one of Big State's students possessed non-purple hair. I might even say the negation of my statement is "All Big State students have non-purple hair."

Phrases such as "all" or "there exists" or "some" are referred to as ***quantifiers***. Sometimes it is useful to distinguish between a **universal quantifier**, which asserts that every object in a given collection possesses a certain property, and an **existential quantifier**, which asserts only that at least one object has that property.

What if we say something like "For all $x$, $x + 1$ is an integer"? As it stands, this is not even a proposition. Since the sentence leaves unclear what, precisely, $x$ is, I can make no determination as to whether the statement is true or false. To make it into a proposition I must stipulate that $x$ is a member of some particular set. The truth or falsity of the resulting proposition would then depend on the set I choose. Since the thing that $x$ represents varies depending on the set of which it is a member, we say that $x$ is a ***variable***.

For example, suppose I assume that $x$ is an integer. In that case I might rephrase the sentence to say, "For all integers $x$, $x + 1$ is an integer," or "For all $x$, if $x$ is an integer then $x + 1$ is also an integer." In this case the proposition is true. But if instead I say that $x$ is a rational number (a fraction), then my sentence becomes, "For all rational numbers $x$, $x+1$ is an integer." This sentence is false. To show that it is false I need only observe that $\frac{1}{2}$ is a rational number but $\frac{1}{2}+1=\frac{3}{2}$ which is not an integer. In other words, I have shown that the statement, "There exists at least one rational number $x$ such that $x+1$ is not an integer," is true. It follows that my original statement is false.

But now suppose that I am less ambitious and say, "There exists a rational number $x$ such that $x+1$ is an integer." That statement is true, since the number $\frac{5}{1}+1=6$ is an integer.

Constructions such as the foregoing are ubiquitous in mathematics. It often happens that we are confronted with a particular set and wish to assert that some or all of the elements of that set have or lack a certain property. Since mathematicians are fond of brevity in their writing, they have developed certain abbreviations for writing such propositions.

For example, instead of writing "For all" it is customary to write an upside-down A, like this: $\forall$. In place of "There exists", we use a backward E, like this: $\exists$. Thus, the statement "For all $x$, if $x$ is an integer then $x+1$ is an integer," would be abbreviated, "$\forall x$, if $x$ is an integer then $x+1$ is an integer." Similarly, we could write "$\exists x$ such that $x$ is a rational number and $x+1$ is an integer."

Let us stipulate that for the rest of this section $x$ will represent an integer. We let $P$ denote some proposition that may or may not be true for a particular integer $x$. For example, $P$ might be the proposition "$x+1$ is an integer." Then $P$ is true regardless of the particular integer $x$ represents. On the other hand, if $P$ is the proposition "$x$ is a perfect square," then $P$ will be true if $x=0,1,4,9,16,25,36,\ldots$ and so on, but it will be false if $x$ is any other integer. It is also possible to have a proposition that is not true for any integer $x$. For example, if $P$ asserts that "$x^2$ is a negative number", then $P$ is false regardless of the value of $x$.

Once we have decided on what, exactly, $P$ is, the assertion that $P$ is true for a given integer $x$ can be abbreviated as $P(x)$, read "$P$ of $x$". Thus, if $P$ is the proposition "$x+1$ is an integer", then the assertion that "For all integers $x$, $x+1$ is an integer" can be abbreviated to $\forall x\ P(x)$. A considerable improvement in efficiency, if not in clarity.

## 2.6 Negations of Quantified Statements

Now suppose that $Q$ is any proposition that may or may not be true for an integer $x$. What can we say about $\neg(\forall x\ Q(x))$? We seek the negation of the statement "For all integers $x$, the proposition $Q$ is true for $x$. We have already seen that the negation of a "for all" statement is a "there exists" statement. Thus, if it is not true that for all integers $x$ the proposition $Q$ is true of $x$, then it must be the case that there exists at least one integer $x$ such that the proposition $Q$ is false for $x$. This can be abbreviated by $\exists x\ \neg(Q(x))$. We could even say that $\neg(\forall x\ Q(x))$ is logically equivalent to $\exists x\ \neg(Q(x))$.

Having absorbed that fact, the reader will be unsurprised to learn that $\neg(\exists x\ Q(x))$ is logically equivalent to $\forall x\ \neg Q(x)$.

Of course, very large books get written on the subtleties of formal logic. Happily, what we have already laid out will be an adequate foundation for what is to come.

## 2.7 Problem Solving Skills

In learning mathematics, understanding all of the terminology and notation is only half the battle. The other half is being able to use this understanding to solve actual problems. In this section we will, by considering several examples, learn some techniques for making problems in logic more approachable.

### 2.7.1 Constructing Truth Tables

How do we go about constructing the truth table for the following proposition:

$$(P \to (Q \vee R)) \to (\neg Q \wedge R)?$$

In approaching any sort of mathematical problem, the first step is be clear on what the problem is asking for. In this case we seek a truth table for a proposition. Thus, I am going to need to consider all possible combinations of truth values for $P$, $Q$, and $R$, and for each one determine whether the proposition above is true or false.

Since there are three propositions, each of which can be either true or false, we conclude there will be eight rows in our table. And now we come to our first dilemma. We need some systematic way of listing all combinations of truth values for the three propositions (for example, we could have $P$ is true, $Q$ is true, and $R$ is false; or we could have all three of them be false; and so on). Otherwise we will end up listing five or six of them, and then struggling to figure out which ones we overlooked. This problem becomes even more acute in cases where we have more than three variables to consider.

Happily, there is a solution to this problem:

| $P$ | $Q$ | $R$ |
|---|---|---|
| T | T | T |
| T | T | F |
| T | F | T |
| T | F | F |
| F | T | T |
| F | T | F |
| F | F | T |
| F | F | F |

Study this table carefully and you will notice that the first column has four T's followed by four F's. The second column has the T's and F's coming in pairs. And the third column alternates T's and F's. This method always works. Look back at the two-variable tables we did earlier. Notice that the first column always has two T's then two F's, while the second column alternates T's and F's. If we had started with four variables, the first column would have had eight T's followed by eight F's. The second column would have four T's, then four F's, then four T's and finally another four F's. Learning this simple trick will spare you a lot of frustration when constructing truth tables.

Now for the proposition itself. We must determine what columns we will need in our table to help us assess its truth value. To do that, we begin breaking the proposition down into its component parts. We see this is an if-then statement whose antecedent is $P \to (Q \vee R)$ and whose conclusion is $\neg Q \wedge R$. Therefore, both of those will be columns in our tables.

But those statements are themselves complex. So I notice that the first part is itself an if-then, this time with $P$ as its antecedent and $Q \vee R$ as its conclusion. The second part is a conjunction with parts $\neg Q$ and $R$.

Now that I have broken apart my proposition, I am ready to construct my table. As I did previously, I will refer to the entire proposition as $Z$.

| $P$ | $Q$ | $R$ | $Q \vee R$ | $\neg Q$ | $P \to (Q \vee R)$ | $\neg Q \wedge R$ | $Z$ |
|---|---|---|---|---|---|---|---|
| T | T | T | T | F | T | F | F |
| T | T | F | T | F | T | F | F |
| T | F | T | T | T | T | T | T |
| T | F | F | F | T | F | F | T |
| F | T | T | T | F | T | F | F |
| F | T | F | T | F | T | F | F |
| F | F | T | T | T | T | T | T |
| F | F | F | F | T | T | F | F |

### 2.7.2 Knight/Knave Problems

As an amusing way of testing your understanding of "The Rules", consider the following scenario.

Let us suppose that we are on an island inhabited entirely by two kinds of people: Knights and Knaves. Knights always tell the truth and Knaves always lie. Suppose that you meet a person who makes the following two statements:

- I like cheese.
- If I like cheese, then I also like pasta.

Can we determine whether this person is a Knight or a Knave?

To answer that, we might reason as follows: He is either a Knight or a Knave. Let us suppose for the moment he is a Knave. In that case everything he says is a lie. So his first statement, that he likes cheese, is false. Now look at the second statement. Since he is a Knave, this statement must be false as well. Since this is an if-then statement, we conclude that the first part is true while the second part is false. But since we have already concluded that the first part is false, we know that actually the conditional statement is true. We would then have a Knave making a true statement, which is impossible. Consequently, this person is a Knight.

There were two main ingredients to solving this problem. The first was a willingness to make an arbitrary guess, just to get the reasoning started. The second involved thinking carefully about when an if-then statement is true, and when it is false.

Here is another example. This time you meet two people who make the following statements:

- A: B is a Knight.
- B: A and I are either both lying or both telling the truth.

What are A and B?

As a place to start, let us suppose that A is a Knave. Then he is lying. It follows that B is not a Knight, and so must be a Knave. Since B is a Knave, we know he is lying as well. So A and B are both lying. But now look at what B actually says. He is making an "or" statement whose first part, that both of them are lying, is true. Consequently, his entire statement is true, which is a contradiction. The only way out of this is to suppose that actually A is a Knight. He would then be telling the truth, which implies that B is also a Knight.

What if we had started instead with the assumption that B is a Knave? Then we would have reasoned that B is lying. In that case we must have that A is telling the truth, for otherwise the "or" statement made by B would be true. But A can't be telling the truth, because he is asserting that B is a Knight, which is contrary to our assumption. This contradiction shows, once again, that A and B must both be Knights.

### 2.7.3 Determining if an Argument is Valid

We now consider one final kind of problem. Consider the following argument:

- $\neg R$
- $P \to R$
- $Q \to R$
- Therefore, $\neg(P \wedge Q)$.

Is this argument valid?

This time we begin by reminding ourselves that a valid argument is one in which the conclusion follows logically from the premises. In other words, if you assume the premises are true, then you find that the conclusion is true as well.

We could make a truth table to help us answer this question, but that seems like rather a lot of work. Instead, let us begin by noticing that since we are assuming that $\neg R$ is true, we must have that $R$ is false. Now look at the remaining two premises. Each of them is a conditional statement with $R$ as a conclusion. Since we know $R$ is false, these statements can be true only if $P$ and $Q$ are both false (remember that a conditional statement with a false antecedent must be true!) Since $P$ and $Q$ are both false, we have that $\neg(P \wedge Q)$ is true. This shows that we have a valid argument.

## 2.8 Problems

The first eight problems take place on the island of Knights and Knaves. Remember that Knights always tell the truth and Knaves always lie.

1. One day you encounter two people who make the following statements:

   A: I am a Knight.

   B: Both of us are lying.

   What can you tell me about A and B?

2. The next day you are investigating the theft of some cookies. You know that one of two people is guilty, and they make the following statements:

   A: I am not the thief.

   B: Exactly one of us is lying.

   Who stole the cookies?

3. The next day you meet two people who make the following statements:

   A: I am a Knight if and only if B is a Knave.

   B: If I am a Knave then A is a Knight.

   What can you tell me about A and B?

4. This time you meet three people. You ask A whether he is a Knight or a Knave. His answer is an incoherent mumble. You then ask B to clarify what A just said. B replies: "A said that he was a Knave." At this point C chimes in by saying "Pay no attention to B, he is lying." What can you tell me about A, B and C?

5. Another time you meet three different people who make the following statements:

   A: If B is a Knight then C is a Knave.

   B: A is a Knave or C is a Knight.

   C: A is a Knave and B is a Knave.

   Again, what can you tell me about A, B and C?

6. This time the three people say:

   A: There is exactly one Knight among us.

   B: There is exactly one Knave among us.

   C: We are all Knights or we are all Knaves.

   What are A, B and C?

7. Once again you find yourself investigating a theft, only this time there are four suspects. They make the following statements:

   A: C is the thief.

   B: I am not the thief.

   C: D is the thief.

   D: C is lying.

   If exactly one of the suspects is a Knight, then who is the thief? Would your answer change if you knew instead that exactly one of the four is a Knave?

8. You encounter six people who make the following statements:

   A: B is a Knave or D is a Knave.

   B: D and F are the same kind.

   C: A is a Knight or F is a Knave.

   D: B is a Knight.

   E: A is a Knight.

   F: B and myself are both Knights or both Knaves.

Who are the Knights and who are the Knaves?

9. Write down the truth tables for the propositions $(P \to Q) \to R$ and $P \to (Q \to R)$. Are these propositions logically equivalent?

10. Write down the truth table for the proposition:
$$(P \wedge Q) \vee (Q \wedge R).$$

11. Write down the truth table for the proposition:
$$[P \vee (Q \wedge R)] \leftrightarrow ([P \vee Q] \wedge [P \vee R]).$$

12. Show that the proposition
$$\neg(P \wedge Q) \leftrightarrow ((\neg P) \vee (\neg Q))$$
is a tautology by writing down its truth table. Verify that it is a tautology by translating it into common English and observing that it says something obvious.

13. Repeat the previous problem, this time with the proposition:
$$[(P \to Q) \vee (P \to R)] \to (\neg(Q \vee R) \to \neg P).$$

14. Find propositions $P$, $Q$ and $R$ such that one of the two compound propositions $(P \leftrightarrow Q) \to R$ and $P \leftrightarrow (Q \to R)$ is true and the other is false.

15. Which of the following are propositions?
    (a) $2^2 + 3^2 = 17$.
    (b) $8x^3 + 6x^2 - 4x + 2$.
    (c) If $n$ is a positive integer, then the sum of the first $n$ positive integers is given by $\frac{n(n+1)}{2}$.
    (d) Will you marry me?
    (e) The king of clubs.

16. Explain how to write the negation of a conditional in the form of a conjunction.

17. Explain how to write the negation of a biconditional in the form of a disjunction.

18. Determine which of the following arguments are valid (note that the commas are used to separate the individual premises in the argument, and that the conclusion comes after the word "therefore"):

    (a) $P \vee Q$, $(\neg Q) \vee R$, $\neg R$ therefore $P$.

    (b) $P \rightarrow Q$, $Q \rightarrow R$, $Q \vee R$ therefore $P$.

    (c) $(\neg P) \vee (\neg Q)$, $R \vee (\neg Q)$, $\neg P$ therefore $R \vee (\neg P)$.

    (d) $P \vee (\neg Q)$, $R \vee (\neg Q)$, $\neg P$ therefore $\neg R$.

19. Let $x$ represent a positive integer, let $S(x)$ be the proposition "$x$ is a perfect square", let $E(x)$ be the proposition "x is even", let $O(x)$ be the proposition "$x$ is odd", and let $C(x)$ be the proposition "$x$ is a perfect cube." Determine the truth value of each of the following propositions. Where appropriate, give an example of a possible value of $x$ that shows the statement to be either true or false. In each case, give a one sentence explanation of your reasoning.

    (a) $S(2)$.

    (b) $S(4) \wedge C(8)$.

    (c) $(E(5) \vee C(27)) \rightarrow (O(11) \wedge S(17))$.

    (d) $\exists x\ (E(x) \wedge O(x))$.

    (e) $\exists x\ (S(x) \wedge C(x))$.

    (f) $\forall x\ (E(x) \vee O(x))$.

    (g) $\forall x\ (S(x) \vee \neg C(x))$.

    (h) $\forall x\ ((E(x) \wedge S(x)) \rightarrow (O(x+1) \wedge \neg(S(x+1))))$.

    (i) $\forall x\ (E(x) \vee E(x+1))$.

    (j) $\exists x\ (S(x) \wedge E(x) \wedge O(\sqrt{x}))$.

20. Let $x$ and $y$ be positive integers and let $P$ and $Q$ be propositions about them. Write down the negation of the following proposition:

$$\forall x\ (\exists y\ (P(x) \wedge Q(y))).$$