From Moore and Parker, CRITICAL THINKING, 9th Ed. (New York: McGraw-Hill, 2009)

that the whole matter of a symbolic system is unfamiliar to you, so we'll start from absolute scratch. Keep in mind, though, that everything builds on what goes before. It's important to master each concept as it's explained and not fall behind. Catching up can be very difficult. If you find yourself having difficulty with a section or a concept, put in some extra effort to master it before moving ahead. It will be worth it in the end.

TRUTH TABLES AND THE TRUTH-FUNCTIONAL SYMBOLS

Our "logical vocabulary" will consist of claim variables and truth-functional symbols. Before we consider the real heart of the subject, truth tables and the symbols that represent them, let's first clarify the use of letters of the alphabet to symbolize terms and claims.

Claim Variables

In Chapter 8, we used uppercase letters to stand for terms in categorical claims. Here, we use uppercase letters to stand for claims. Our main interest is now in the way that words such as "not," "and," "or," and so on affect claims and link them together to produce compound claims out of simpler ones. So, don't confuse the Ps and Qs, called **claim variables**, that appear in this chapter with the variables used for terms in Chapter 8.*

Truth Tables

Let's now consider truth tables and symbols. In truth-functional logic, any given claim, P, is either true or false. The following little table, called a **truth table**, displays both possible truth values for P:

P
T
F

Whichever truth value the claim P might have, its negation or contradictory, which we'll symbolize ~P, will have the other. Here, then, is the truth table for ***negation:***

P	~P
T	F
F	T

The left-hand column of this table sets out both possible truth values for P, and the right-hand column sets out the truth values for ~P based on P's values. This is a way of defining the negation sign, ~, in front of the P. The symbol means "change the truth value from T to F or from F to T, depending on

* It is customary to use one kind of symbol, usually lowercase letters or Greek letters, as *claim variables* and plain or italicized uppercase letters for *specific claims.* Although this use has some technical advantages and makes possible a certain theoretical neatness, students often find it confusing. Therefore, we'll use uppercase letters for both variables and specific claims and simply make it clear which way we're using the letters.

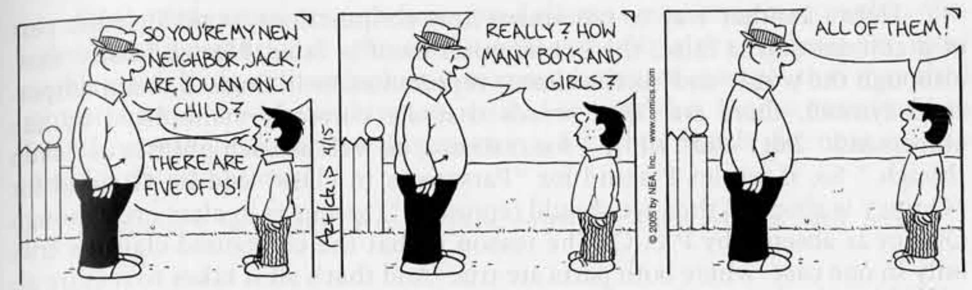

■ The word "and," when used in questions, can produce some interesting and amusing results. In this case, Brutus means to ask, "How many of them are boys, and how many of them are girls?" But Jack thinks he asks, "How many of them are girls or boys?" There's even a third version: "How many of them are *both* girls and boys?" Presumably, none.

P's values." Because it's handy to have a name for negations that you can say aloud, we read ~P as "not-P." So, if P were "Parker is at home," then ~P would be "It is not the case that Parker is at home," or, more simply, "Parker is not at home." In a moment we'll define other symbols by means of truth tables, so make sure you understand how this one works.

Because any given claim is either true or false, two claims, P and Q, must both be true, both be false, or have opposite truth values, for a total of four possible combinations. Here are the possibilities in truth-table form:

P	Q
T	T
T	F
F	T
F	F

A **conjunction** is a compound claim made from two simpler claims, called *conjuncts. A conjunction is true if and only if both of the simpler claims that make it up (its conjuncts) are true.* An example of a conjunction is the claim "Parker is at home and Moore is at work." We'll express the conjunction of P and Q by connecting them with an ampersand (&). The truth table for conjunctions looks like this:

P	Q	P & Q
T	T	T
T	F	F
F	T	F
F	F	F

P & Q is true in the first row only, where both P and Q are true. Notice that the "truth conditions" in this row match those required in the italicized statement above.*

* Some of the words that have truth-functional meaning have other kinds of meanings as well. For example, "and" can signify not only that two things happened but that one happened earlier than the other. An example: "Melinda got on the train and bought her ticket" is quite different from "Melinda bought her ticket and got on the train." In this case, "and" operates as if it were "and then."

Here's another way to remember how conjunctions work: If either part of a conjunction is false, the conjunction itself is false. Notice finally that, although the word "and" is the closest representative in English to our ampersand symbol, there are other words that are correctly symbolized by the ampersand: "but" and "while," for instance, as well as such phrases as "even though." So, if we let P stand for "Parsons is in class" and let Q stand for "Quincy is absent," then we should represent "Parsons is in class even though Quincy is absent" by P & Q. The reason is that the compound claim is true only in one case: where both parts are true. And that's all it takes to require an ampersand to represent the connecting word or phrase.

A **disjunction** is another compound claim made up of two simpler claims, called *disjuncts. A disjunction is false if and only if both of its disjuncts are false.* Here's an example of a disjunction: "Either Parker is at home, or Moore is at work." We'll use the symbol ∨ ("wedge") to represent disjunction when we symbolize claims—as indicated in the example, the closest word in English to this symbol is "or." The truth table for disjunctions is this:

P	Q	P ∨ Q
T	T	T
T	F	T
F	T	T
F	F	F

Notice here that a disjunction is false only in the last row, where both of its disjuncts are false. In all other cases, a disjunction is true.

The third kind of compound claim made from two simpler claims is the **conditional claim.** In ordinary English, the most common way of stating conditionals is by means of the words "if . . . then . . . ," as in the example "If Parker is at home, then Moore is at work."

We'll use an arrow to symbolize conditionals: P → Q. The first claim in a conditional, the P in the symbolization, is the **antecedent,** and the second—Q in this case—is the **consequent.** *A conditional claim is false if and only if its antecedent is true and its consequent is false.* The truth table for conditionals looks like this:

P	Q	P → Q
T	T	T
T	F	F
F	T	T
F	F	T

Only in the second row, where the antecedent P is true and the consequent Q is false, does the conditional turn out to be false. In all other cases, it is true.*

*Like the conjunction, conditionals in ordinary language can have more than the meaning we assign to the arrow. The arrow represents what is often called the "material conditional," conditionals that are true except when the antecedent is true and the consequent false.

Differences between material conditionals and the conditionals used in ordinary language have held the attention of logicians and philosophers for a long time and are still controversial. See, for example, Richard Bradley, "A Defence of the Ramsey Test," in the January 2007 issue of the philosophical journal *Mind* (Vol. 116, Number 461, pp. 1-21).

Of the four types of truth-functional claims—negation, conjunction, disjunction, and conditional—the conditional typically gives students the most trouble. Let's have a closer look at it by considering an example that may shed light on how and why conditionals work. Let's say that Moore promises you that, if his paycheck arrives this morning, he'll buy lunch. So, now we can consider the conditional

> If Moore's paycheck arrives this morning, then Moore will buy lunch.

We can symbolize this using P (for the claim about the paycheck) and L (for the claim about lunch): P → L. Now let's try to see why the truth table above fits this claim.

The easiest way to see this is by asking yourself what it would take for Moore to break his promise. A moment's thought should make this clear: Two things have to happen before we can say that Moore has fibbed to you. The first is that his paycheck must arrive this morning. (After all, he didn't say what he was going to do if his paycheck *didn't* arrive, did he?) Then, it being true that his paycheck arrives, he must then *not* buy you lunch. Together, these two items make it clear that Moore's original promise was false. Notice: Under no other circumstances would we say that Moore broke his promise. And *that* is why the truth table has a conditional false in one and only one case, namely, where the antecedent is true and the consequent is false. Basic information about all four symbols is summarized in Figure 1.

FIGURE 1 The Four Basic Truth-Functional Symbols

Negation (~)

Truth table:

P	~P
T	F
F	T

Closest English counterparts: "not," or "it is not the case that"

Conjunction (&)

Truth table:

P	Q	(P & Q)
T	T	T
T	F	F
F	T	F
F	F	F

Closest English counterparts: "and," "but," "while"

Disjunction (∨)

Truth table:

P	Q	(P ∨ Q)
T	T	T
T	F	T
F	T	T
F	F	F

Closest English counterparts: "or," "unless"

Conditional (→)

Truth table:

P	Q	(P → Q)
T	T	T
T	F	F
F	T	T
F	F	T

Closest English counterparts: "if . . . then," "provided that"

Our truth-functional symbols can work in combination. Consider, for example, the claim "If Paula doesn't go to work, then Quincy will have to work a double shift." We'll represent the two simple claims in the obvious way, as follows:

P = Paula goes to work.
Q = Quincy has to work a double shift.

And we can symbolize the entire claim like this:

~P → Q

Here is a truth table for this symbolization:

P	Q	~P	~P → Q
T	T	F	T
T	F	F	T
F	T	T	T
F	F	T	F

Notice that the symbolized claim ~P → Q is false in the *last* row of this table. That's because, here and only here, the antecedent, ~P, is true and its consequent, Q, is false. Notice that we work from the simplest parts to the most complex: The truth value of P in a given row determines the truth value of ~P, and that truth value in turn, along with the one for Q, determines the truth value of ~P → Q.

Consider another combination: "If Paula goes to work, then Quincy and Rogers will get a day off." This claim is symbolized this way:

P → (Q & R)

This symbolization requires parentheses in order to prevent confusion with (P → Q) & R, which symbolizes a different claim and has a different truth table. Our claim is a conditional with a conjunction for a consequent, whereas (P → Q) & R is a conjunction with a conditional as one of the conjuncts. The parentheses are what make this clear.

You need to know a few principles to produce the truth table for the symbolized claim P → (Q & R). First, you have to know how to set up all the possible combinations of true and false for the three simple claims P, Q, and R. In claims with only one letter, there were two possibilities, T and F. In claims with two letters, there were four possibilities. *Every time we add another letter, the number of possible combinations of T and F doubles, and so, therefore, does the number of rows in our truth table.* The formula for determining the number of rows in a truth table for a compound claim is $r = 2^n$, where r is the number of rows in the table and n is the number of letters in the symbolization. Because the claim we are interested in has three letters, our truth

table will have eight rows, one for each possible combination of T and F for P, Q, and R. Here's how we do it:

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

The systematic way to construct such a table is to alternate Ts and Fs in the right-hand column, then alternate *pairs* of Ts and *pairs* of Fs in the next column to the left, then sets of *four* Ts and sets of *four* Fs in the next, and so forth. The leftmost column will always wind up being half Ts and half Fs.

The second thing we have to know is that the truth value of a compound claim in any particular case (i.e., any row of its truth table) depends entirely upon the truth values of its parts; and if these parts are themselves compound, their truth values depend upon those of their parts; and so on, until we get down to letters standing alone. The columns under the letters, which you have just learned to construct, will then tell us what we need to know. Let's build a truth table for P → (Q & R) and see how this works.

In Depth

Test Yourself

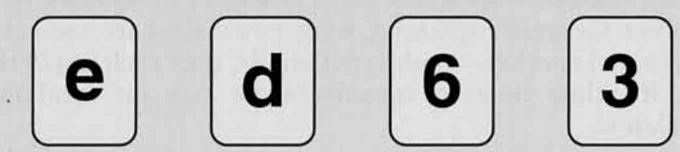

These cards are from a deck that has letters on one side and numbers on the other. They are supposed to obey the following rule: "If there is a vowel on one side, then the card has an even number on the other side."

Question: To see that the rule has been kept, how many cards must be turned over and checked? (Most university students flunk this simple test of critical thinking.)

P	Q	R	Q & R	$P \rightarrow (Q \& R)$
T	T	T	T	T
T	T	F	F	F
T	F	T	F	F
T	F	F	F	F
F	T	T	T	T
F	T	F	F	T
F	F	T	F	T
F	F	F	F	T

The three columns at the left, under P, Q, and R, are our *reference columns,* set up just as we discussed above. They determine what goes on in the rest of the table. From the second and third columns, under the Q and the R, we can fill in the column under Q & R. Notice that this column contains a T only in the first and fifth rows, where both Q and R are true. Next, from the column under the P and the one under Q & R, we can fill in the last column, which is the one for the entire symbolized claim. It contains Fs only in rows two, three, and four, which are the only ones where its antecedent is true and its consequent is false.

What our table gives us is a *truth-functional analysis* of our original claim. Such an analysis displays the compound claim's truth value, based on the truth values of its simpler parts.

If you've followed everything so far without problems, that's great. If you've not yet understood the basic truth table idea, however, as well as the truth tables for the truth-functional symbols, then by all means stop now and go back over this material. You should also understand how to build a truth table for symbolizations consisting of three or more letters. What comes later builds on this foundation, and as with any construction project, without a strong foundation the whole thing collapses.

A final note before we move on: Two claims are **truth-functionally equivalent** if they have exactly the same truth table—that is, if the Ts and Fs in the column under one claim are in the same arrangement as those in the column under the other. Generally speaking, when two claims are equivalent, one can be used in place of another—truth-functionally, they each imply the other.

Okay. It's time now to consider some tips for symbolizing truth-functional claims.

SYMBOLIZING COMPOUND CLAIMS

Most of the things we can do with symbolized claims are pretty straightforward; that is, if you learn the techniques, you can apply them in a relatively clear-cut way. What's less clear-cut is how to symbolize a claim in the first place. We'll cover a few tips for symbolization in this section and then give you a chance to practice with some exercises.

Remember, when you symbolize a claim, you're displaying its truth-functional structure. The idea is to produce a version that will be truth-functionally equivalent to the original informal claim—that is, one that will be true under all the same circumstances as the original and false under all

In Depth

Truth-Functional Logic and Electrical Circuits

We mentioned at the beginning of the chapter that truth-functional logic is the basis of digital computing. This is because, translated into hardware systems, "true" and "false" become "on" and "off." Although there's a lot more to it than this, we can illustrate in a crude way a little of how this works.

Let's construct a simple electrical circuit from an electrical source to a ground and put a lightbulb in it somewhere, like this:

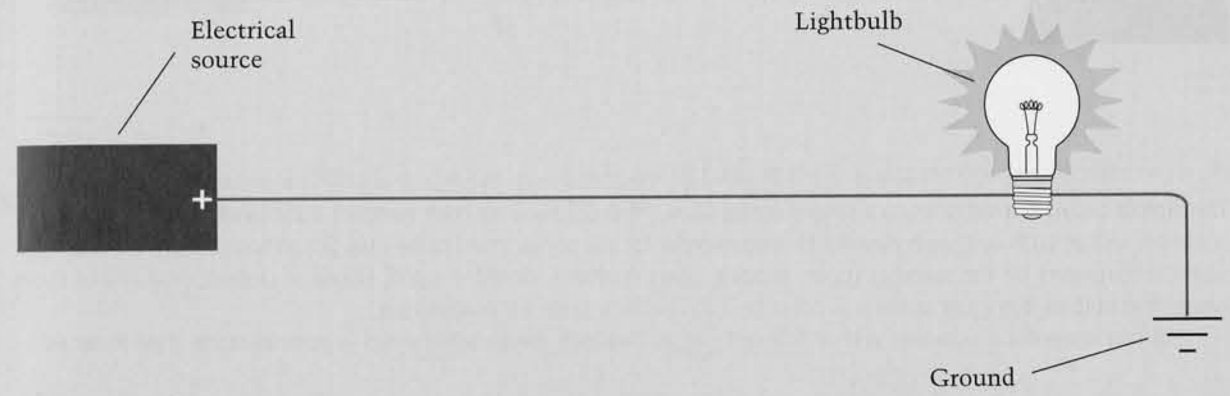

In this situation, the light burns all the time. Now, let's add a switch and give it a name, "P," like so:

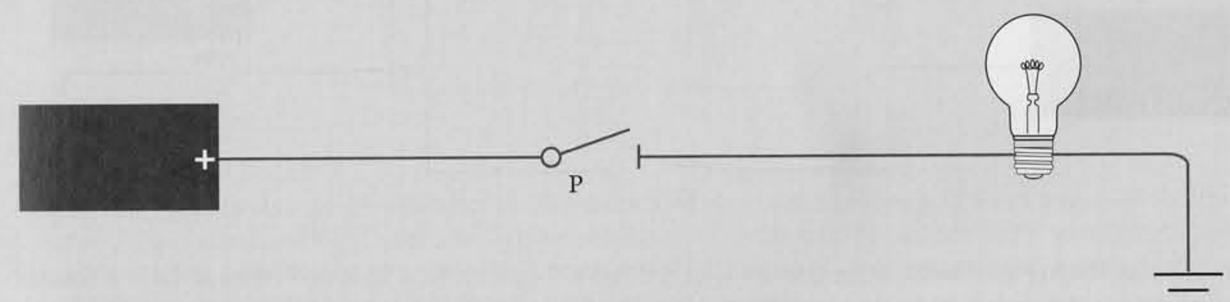

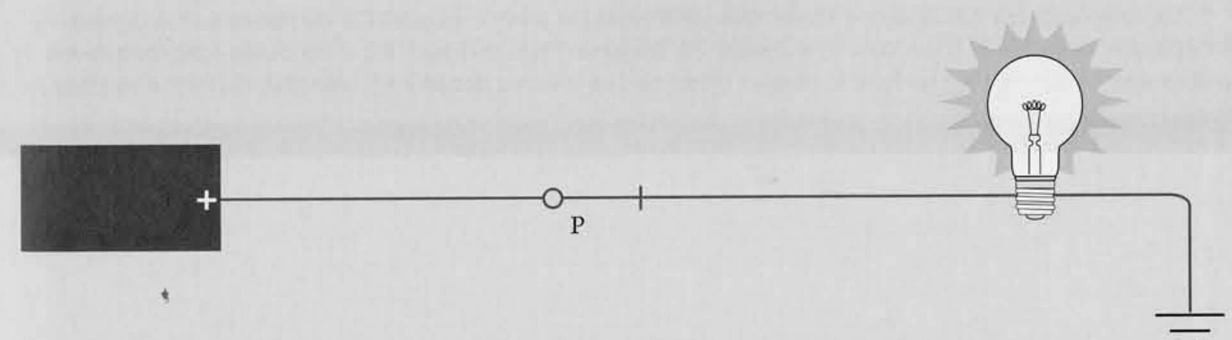

that Q must be true to preserve the truth of the first premise. And that completes the assignment:

P	Q	R	S	T
T	T	F	F	T

This is one row in the truth table for this argument—the only row, as it turned out—in which all the premises are true and the conclusion is false; thus, it is the row that proves the argument invalid.

In the preceding example, there was a premise that forced us to begin with a particular assignment to a letter. Sometimes, neither the conclusion nor any of the premises forces an assignment on us. In that case, we must use trial and error: Begin with one assignment that makes the conclusion false (or some premise true) and see if it will work. If not, try another assignment. If all fail, then the argument is valid.

In Depth

Common Truth-Functional Argument Patterns

Some truth-functional patterns are so built into our thinking process that they almost operate at a subverbal level. But, rather than trust our subverbal skills, whatever those might be, let's identify three common patterns that are perfectly valid—their conclusions follow with certainty from their premises—and three invalid imposters—each of the imposters bears a resemblance to one of the good guys. We'll set them up in pairs:

Valid Argument Forms

In these cases, the premises guarantee the conclusion.

1. Modus ponens (or affirming the antecedent)
   If P, then Q
   P
   ———
   Q
2. Modus tollens (or denying the consequent)
   If P, then Q
   Not-Q
   ———
   Not-P
3. Chain argument
   If P, then Q
   If Q, then R
   ———
   If P, then R

Invalid Argument Forms

Here, the premises can be true while the conclusion is false.

1-A. Affirming the consequent
   If P, then Q
   Q
   ———
   P

2-A. Denying the antecedent
   If P, then Q
   Not-P
   ———
   Not-Q

3-A. Undistributed middle (truth-functional version)
   If P, then Q
   If R, then Q
   ———
   If P, then R

Often, several rows of a truth table will make the premises true and the conclusion false; any one of them is all it takes to prove invalidity. Don't get the mistaken idea that, just because the premises are all true in one row and so is the conclusion, the conclusion follows from the premises—that is, that the argument must be valid. To be valid, the conclusion must be true in *every* row in which all the premises are true.

To review: Try to assign Ts and Fs to the letters in the symbolization so that all premises come out true and the conclusion comes out false. There may be more than one way to do it; any of them will do to prove the argument invalid. If it is impossible to make the premises and conclusion come out this way, the argument is valid.

Exercise 9-4

Construct full truth tables or use the short truth-table method to determine which of the following arguments are valid.

▲ **1.** P ∨ ~Q
   ~Q
   ———
   ~P

2. P → Q
   ~Q
   ———
   ~P

3. ~(P ∨ Q)
   R → P
   ———
   ~R

▲ **4.** P → (Q → R)
   ~(P → Q)
   ———
   R

5. P ∨ (Q → R)
   Q & ~R
   ———
   ~P

6. (P → Q) ∨ (R → Q)
   P & (~P → ~R)
   ———
   Q

▲ **7.** (P & R) → Q
   ~Q
   ———
   ~P

8. P & (~Q → ~P)
   R → ~Q
   ———
   ~R

9. L ∨ ~J
   R → J
   ———
   L → ~R