A BRIEF INTRODUCTION TO LOGIC FOR METAPHYSICIANS

0. Logic, Probability, and Formal Structure

Logic is often divided into two distinct areas, *inductive* logic and *deductive* logic. *Inductive* logic is concerned with probabilities. That is, in an inductive argument we attempt to show that if the premises of the argument are (all) true, then the conclusion is *probably* true. Consider the following argument:

Most men are bachelors.
Jones is a man.
So, Jones is (probably) a bachelor.

If the premises of this argument are true, then the conclusion is probably true. But, of course, it is *possible* that the conclusion is false, even given the truth of the premises. This kind of argument only *attempts* to establish probability. We will, in general, not be concerned with this kind of argument.

Deductive arguments, on the other, attempt to show that if the premises are true, then the conclusion *must* be true—i.e., that it *cannot possibly* be false. It is deductive arguments that we will be concerned with in this introduction to logic.

An important feature of deductive arguments is that their success or failure depends upon the *form* of the argument, not its specific *content*. Consider the following argument:

All bachelors are male.
Jones is a bachelor.
So, Jones is male.

In this argument, the truth of the premises guarantees the truth of the conclusion. But note that this has nothing to do, specifically, with Jones, bachelorhood, or with being male. The argument is just an instance of the following form:

All A's are B's
x is an A.
So, x is a B.

It doesn't matter what we "plug in" for "A," "B" or "x." And this means that *any* argument with this form will always have the property of being such that if the premises are true, then the conclusion must be true.

The point here is that in deductive reasoning, we are concerned with the form of arguments. But arguments are composed of statements (or "*propositions*"), and this means that we must be concerned with the form or internal structure of statements. The following sections discuss methods for highlighting differing degrees of the internal structure of statements.

1. Truth functional logic (i.e., propositional logic)

Since logic is concerned with the evaluation of arguments, and arguments are composed of statements (i.e., *propositions*), logic must be concerned with the structure of statements. As we shall see, we can, depending upon how we represent a given statement, exhibit more or less of its structure. Consider the following argument:

> If Bush is a Republican, he voted for Cheney.
> Bush is a Republican.
> So, Bush voted for Cheney.

This argument is composed of three statements, two premises, and the conclusion. Suppose we represented each distinct statement with a distinct symbol—let us use capital letters. Then, we would represent the argument as:

> A
> B
> $\therefore$C (“$\therefore$” = “therefore”)

As we can see, this way of representing the statements that comprise this argument show us nothing interesting about it. These three symbols could just as well represent any three statements, whether or not they constituted a good argument. In this case, our way of representing the statements shows us nothing about their internal structure. But it is an immediately obvious feature of the first statement of our argument that it in some sense “contains” both of the other statements. We need some way to represent this.

To do this, we will introduce “truth functional logic.” Truth functional logic is the logic of statements where the truth value (i.e., truth or falsity) of statements is a *function* of its components. Consider a simple example:

> It is not the case that Bush is President.

The truth value of this statement is a function of its component,

> Bush is President.

What this means is that the truth or falsity of the first statement is entirely determined by (in mathematical terms, *is a function of*) the truth value of the second. In general, if I know the truth value of “P,” I also know (or can infer) the truth value of “It is not the case that P.” We will say that a statement is “truth functionally atomic” if its truth value is not a function of any of its “proper” (truth functional) components. (A “proper” component of a statement is any “smaller” statement contained in it—i.e., any “part” of it *other than itself*.) We will also say that a statement is “truth functionally compound” if it is not truth functionally atomic. (Note that not all compound statements are *truth functionally compound*. “Cheney believes that Bush is President” contains “Bush is President as a (*non*-truth functional) part, but the truth or falsity of the part does not in itself *determine* the truth or falsity of the larger compound.) Until we complicate matters in the next section, we will use capital letters to stand for truth functionally atomic statements.

Although not all are strictly necessary, there are typically five ways of forming truth functional compounds, that is, five kinds of truth functional compound statements: negations ("not" statements), conjunctions ("and" statements"), disjunctions ("or" statements), conditionals ("if … then …" statements), and biconditionals ("if and only if" statements). How these compounds function is usually given by truth tables. In what follows, let us use lower case letters to stand for "*propositional variables*"—i.e., "p" and "q" can stand for any statements at all, compound or atomic.

Negations are easy. The truth value of the negation of a statement is simply the opposite of the truth value of the original statement. The sign we will use to represent negation is "tilde"—"~". The truth table for negation is:

p	~p
T	F
F	T

Each row in the table represents a distinct distribution of truth values to the components of the compound in question—in the case, a negation. Since there is only one (proper) component here, there are two possibilities. (In general, where there are *n* distinct components, there are 2^n distinct distributions of truth values to those components.) So the table here says that for any statement "p," if "p" is true, "~p" is false, and if "p" is false, then "~p" is true. Let us construct truth tables for the other compounds in like fashion.

The symbol for conjunction is the *dot*—"·", and the *wedge*—"v" for disjunction. (The specific symbols used often vary.) The tables are as follows:

p	q	p · q
T	T	T
T	F	F
F	T	F
F	F	F

p	q	p v q
T	T	T
T	F	T
F	T	T
F	F	F

In other words, a conjunction is true just in case both its "*parts*" (its "*conjuncts*") are true, and a disjunction is false just in case both its parts (its "*disjuncts*") are false.

The symbol for conditionals is the *horseshoe*—"⊃", and the *triple bar*—"≡" for biconditionals. Their truth tables are as follows:

p	q	p ⊃ q
T	T	T
T	F	F
F	T	T
F	F	T

p	q	p ≡ q
T	T	T
T	F	F
F	T	F
F	F	T

Note that a conditional is false just in case its "antecedent" (what comes before the horseshoe—the "if" clause) is true *and* its "consequent" (what comes after the horseshoe—the "then" clause) is false. This may seem odd, in that not all English statements have this feature. But at least *some* do (and the rest are not *truth functional conditionals*), and that is how we will treat them for now. Biconditionals are true just in case both their components have the same truth value, i.e., either are both true or are both false.

Now we are in a position to better represent the argument we consider earlier. Let "R" stand for "Bush is a Republican," and "C" stand for "Bush voted for Cheney." Now the argument comes out:

$R \supset C$
R
$\therefore C$

And now we can see why this is a good argument. Assuming that the premises are true, we know that both $R \supset C$ and R are true. But then, given the truth table for "$\supset$," we know that C must be true. (That is, for every line on the truth table where both the conditional and its antecedent are true, the consequent is also true. See for yourself.) So, we know that *if* the premises are true, then the conclusion *must* be true—it cannot possibly be false.

This, in turn, gives us the definition of a *valid* argument: An argument is *valid* if the truth of the premises guarantees the truth of the conclusion. *In other words,* ***if*** *the premises are all true, the conclusion* ***cannot possibly be false—****it* ***must be*** *true*.

Note that an argument can be valid even if one or more of its premises are false, or even if its conclusion is false. Consider:

If Cheney is a Democrat, he voted for Jay Leno.
Cheney is a Democrat.
Therefore, Cheney voted for Jay Leno.

Again, *were* the premises true (even though they *probably* aren't!) then the conclusion would have to be true. So validity is an "iffy" property of arguments. To say that an argument is valid is **not** to say that any of the statements that make it up are in fact true. It is only to say that **if** the premises are all true, then the conclusion **must** be true.

Finally, we must occasionally use parentheses in our symbolic representations of statements in order to make clear the "scope" of each truth functional connective. Consider:

Cheney is a Democrat or Bush is a Republican and Bush is President.

How is this to be understood? Does it say 1) that *either* Cheney is a Democrat *or* it is true both that Bush is a Republican and that Bush is President? Or does it say 2) that *both* of the statements "Either Cheney is a Democrat or Bush is a Republican" *and* "Bush is President" are true? That is, is this a disjunction whose second disjunct is a conjunction, or is it a conjunction whose first conjunct is a disjunction? To distinguish these possibilities, we need parentheses. The first possibility would be represented:

$$C \vee (R \cdot P)$$

while the second comes out:

$$(C \vee R) \cdot P.$$

So much for truth functional logic.

2. Predicate Logic

We said earlier that logic was concerned with the structure of statements, and that we could, depending upon how we represented them, exhibit more of less of that structure. By assigning simple symbols (in this case, capital letters) to truth functionally atomic statements, we exhibit none of their internal structure. Nevertheless, statements that are truth functionally atomic sometimes contain additional structure that is relevant to evaluating arguments. Consider the following three statements:

Clinton is president.
Bush is president.
Bush is married.

You will note that the middle statement has something (different) in common with each of the other two statements. But since each is truth functionally atomic, each would be represented by a distinct capital letter, which would exhibit none of this internal structural similarity.

Consequently, we may want to complicate the way we represent truth functionally atomic statements. Intuitively, each of these statements is composed of a subject and a predicate. The first two share the same predicate ("… is president"), while the last two share the same subject ("Bush"). To exhibit this, we will complicate our language in the following way: we will use capital letters to stand for predicates, and lower case letters to stand for names, or other referring expressions. So let "P" stand for "… is president," "M" stand for "… is married," "c" stand for "Clinton," and "b" stand for "Bush." (There is a general convention to use letters at the beginning of the alphabet for names.) Now we can make the structural similarity of these three statements apparent by representing them thus:

Pc Pb Mb.

(We might read the first as "P of c," or "P is true of c," or, more intuitively, "c has property P.")

But of course, not all truth functionally atomic statements are quite so simple. Each of the above attributes a property to a single individual. But we might also want to say that two or more individuals stand in some relation to one another. Consider, "Reagan was president before Bush." Letting "P" stand for "…was president before…," "a" stand for Reagan, and "b" stand for Bush, we could represent this as:

Pab. (This is sometimes written "aPb.")

(That is, "*a* stands in the relation *P* to *b*.")

Of course, there are three place relations ("John gave a gift to Wanda"), four place relations ("Betty gave Bruce an STD which she received from Bob who got it from Bill"), and so on. We call the predicates that stand for properties (i.e., that can be attributed only to one thing at a time) "one place predicates," predicates that stand for relations between two things, "two place predicates," and so on.

At this point, these truth functionally atomic statements can be compounded just as before, and none of this will alter our evaluation of arguments. So, letting "R" stand for "…is a republican," "V" stand for "…voted for …," and "C" stand for "Cheney," we can represent our first argument as:

$Rb \supset Vbc$
Rb
$\therefore Vbc$

Note that within any given example, the "meaning" of any of these letters is clear. That is, each names some specific thing, property, or relation. Thus, each is called a "term": predicates will be called "general terms" and names (or other expressions taken as referring to some specific individual), "individual terms." These are to be distinguished from "variables," which will be introduced in the following section.

So much for predicate logic.

3. Quantificational Logic

Consider the following argument:

All men are mortal.
Socrates is a man.
Therefore, Socrates is mortal.

This argument seems intuitively valid, yet we cannot represent it as such given the tools available to us at this point. Given our understanding of predicate logic, we can see that there is plenty of internal structure here, and we can even see how to represent the last two statements. Letting "a" stand for "Socrates," "H" for "…is a man" (i.e., "…is *human*"), and "M" for "…is mortal," the last two statements become "Ha" and "Ma" respectively."

But how do we represent the first statement? It will help us to "translate" the first statement to the following: Take anything you like, if it is a man, then it is mortal. (In other words, "Everything that is a man is also mortal.") Now, consider the latter part of this "translation," the part that comes after "Take anything you like …." Note that here we find our predicates "…is a man" and "…is mortal" being predicated of "*it*." Note further that *it*, in isolation, has an indeterminate reference. I.e., what "it" stands for is determined by the first part of the sentence, and if you isolate it from that part of the sentence, what is left is nothing but a kind of empty "place holder" for a name or referring expression to be supplied by the general context. "It" in this case functions much like a "variable" in mathematics, and so, as in mathematics, we will use lower case letters towards the end of the alphabet (typically, "x," "y," and "z") as standing for different occurrences of the word "it." (We need more than one variable letter as some sentences contain multiple uses of the word "it" with distinct references.) So, "it is mortal" would be written "Mx" and "it is a man," "Hx." And so,

if it is man then it is mortal

can be represented

$$Hx \supset Mx.$$

We are still not finished, however. What should we do the phrase, "Take anything you like …"? Without this phrase (or something like it), our statement doesn't really make any sense, because we do not know what the "it" is that we are speaking about. So this phrase establishes the reference(s) of "it" in our sentence, and it says the "it" we are speaking about is *any* "it." We are speaking of *all* "its." So, we are in effect saying, "Take anything you like, call it x, if x is a man, then x is mortal." Alternately, "For all x, if x is a man, then x is mortal." So we need a way of representing the phrase "For all x." For this we will introduce what is called the "Universal Quantifier," represented by an upside down "A" ("$\forall$"), followed by the variable to stand for "*it*." So, "All men are mortal" (i.e., "Take anything you like, if it is man, then it is mortal"), can be represented

$$(\forall x)(Hx \supset Mx).$$

Note: in some texts, the upside down "A" is eliminated, and so "(x)" would represent the expression "For all x." Also, in some texts the parentheses around the quantifier "$\forall x$" are omitted. But the other parentheses are more important. They define the "scope" of the quantifier that precedes them. Recall that it is the quantifier that tells us what "it" ("x") stands for (or "ranges over"). But some statements contain more than one quantifier (i.e., multiple uses of the word "it" with distinct references), and so the use of parentheses is then necessary to distinguish the scope of each quantifier (i.e., which "it" it "ranges over" or "binds").

Although not strictly necessary, it is customary to define one additional quantifier. While the universal quantifier stands for "everything," it is often convenient to have another to stand for "something." This is called the "Existential Quantifier," and is represented by a backwards "E" ("$\exists$"). So if I want to say "Something is a man" (i.e., "there is something such that it is man"), I would represent it as

$(\exists x)Hx$.

To see how all of this works, let me "translate" a few statements into our symbolic language with quantifiers. (I will leave it to you to figure out what capital letters I used to represent various predicates.)

Everything is attracted by everything.

$(\forall x)(\forall y)(Axy)$

(Take anything you like, call it x, and anything you like, call it y, x attracts y.)

All bodies attract one another.

$(\forall x)(\forall y)((Bx \cdot By) \supset Axy)$

(Take anything you like, call it x, and anything you like, call it y, if it is true both that x and y are bodies, then x attracts y.)

Everybody loves somebody.

$(\forall x)(Px \supset (\exists y)(Py \cdot Lxy))$

(Take anything you like, call it x, if it (x) is a person then there is something, call it y, such that it (y) is a person, and such that x likes y.)

Somebody loves everybody.

$(\exists x)(Px \cdot (\forall y)(Py \supset Lxy))$

(There is something, call it x, such that x is a person and such that for anything you like, call it y, if y is a person, then x loves y.)

Everybody doesn't love something, but nobody doesn't love Sara Lee.

$(\forall x)(Px \supset (\exists y)(\sim Lxy)) \cdot \sim(\exists x)(Px \cdot \sim Lxs)$

(Take anything you like, call it x, if x is person then there is some thing, call it y, such that it is not the case that x likes y; AND it is not the case that there is a thing, call it x, such that x is a person and does not like Sara Lee.)

Some additional terminology: a variable that occurs within the scope of a quantifier that quantifies over it is called a "bound" variable. (E.g., "x" but not "y" in $(\forall x)Pxy$.) If it is not bound, it is "free." Sentences with "free" variables are sometimes called "open sentences." (An example occurred within the previous parentheses. They are called "open" because they are incomplete. It is as if they had a "hole" in them. Consider "Sx." This simply says that "it is S," but it tells us nothing about what "it" is. "It" is a kind of "place holder" for some specific referring expression, and so in a sense, "Sx" remains "open," or incomplete until it is "closed" by specifying the range of "x" with a quantifier.)

Finally, note that in order to evaluate arguments in which such sentences occur, we will need to formalize some "rules" clarifying the implications of such quantified sentences. We need not go into all of them here, but two are worth mentioning. The first is what is sometimes called "Universal Instantiation." The idea is that if everything has some given property, then any specific thing we name must have that property. The second one is sometimes called "Existential Generalization." This principle says that if we know that some specific thing has some property, we can infer the more general statement that *something* (or other) has this property. So this principle would sanction our reasoning from "Sa" (i.e., the individual *a* has the property *S*) to $(\exists x)Sx$. This second principle turns out to be more controversial than the first.

4. What's in a Name?

A lot, as it turns out. (This section can be read as an introduction to Quine's "On What There Is.")

It is noteworthy that a *lot* of English can be more or less adequately "translated" into the quantified symbolic logic we created above. All of mathematics can be translated into it, as well as the physical sciences. (There are problems when we come to the statements of psychology. Should this be surprising?) So, all (or almost all) of our scientific theories can be translated into this language.

Next, note the connection between our theories or explanations and our ontological commitments. ("*Ontological commitments*" are the implications regarding existence that are found in our beliefs, or in the things we say.) We believe (if we believe) in the existence of quarks because we believe (if we believe) that our best physical theories talk about such things. So (the statements that make up) our explanations of things dictate our ontological commitments. That is, our *best* explanations. This is not to say that *what there is* (really) is determined by the kinds of theories we have, but only that what we can (and must) reasonably *believe* there to be is determined by our explanations. So our ultimate ontological commitments can, with a little work, be "read" from how we talk about things in our best explanations.

But consider the following sentences:

Bob doesn't ski.

and

Santa Claus doesn't exist.

The two statements apparently have the same internal structure. They each apparently say of some object that it fails to have some property. We might represent the structure of the first statement by paraphrasing it as:

There is something, named “Bob,” and it is false that this thing skis.

i.e., $(\exists x)(x \text{ is named “Bob”} \cdot \sim x \text{ skis})$

But if each of these statements has the same internal structure, then the second one should be paraphrased as:

There is something, named “Santa Claus,” and it is false that this thing exists.

i.e., $(\exists x)(x \text{ is named “Santa Claus”} \cdot \sim x \text{ exists})$

But now look what has happened: we have apparently committed ourselves to the existence of Santa Claus in the very statement in which we set out to deny it. We have said that “there is” something, and that this “something” has (or, in this case, fails to have) a certain property. The point here is that we might intuitively look to the *names* (or other referring expressions) in our statements as expressing our ontological commitments. The problem is that we have names for things that do not exist. Are we thus committed to the existence of Santa Claus, flying horses, and Sherlock Holmes?

Question: *What’s a (metaphysically inclined) logician to do?*

Answer: Either change our logic (Quine), or change our metaphysics (Parsons).

So much for our introduction to logic. Now on to Metaphysics!