

Some Notes on Logic

Craig Fox

Summer, 2002; Rev. Fall, 2004

1 Introduction.

1.1 What is our approach to logic?

There are many different ways one might approach the study of logic. For example, we might think of logic as describing “how we think,” or better, “how we *ought* to think.” This way of approaching the subject, however, lends itself to some difficulties. Perhaps the most pressing difficulty might be exactly how we would find out *how we think* (or *ought to*). Would we somehow study people while they think? How would we determine exactly what it was that they were really thinking? How would we determine if someone was thinking “poorly” or “well?”

It does not seem that there are clear answers to these questions. If we proceed in this way, logic is in danger of becoming too *psychological*, we might say. There are a number of reasons why we should be wary of such an outcome. A different approach would probably be better.

The ultimate question with which we will be concerned, with our study of logic, is whether or not some given argument is good or bad. Partially in order to answer questions of this sort in a *general* way, our logic will be a symbolic one. This means that one of our tasks, when trying to assess an argument, will be to “translate” an argument into symbolic notation. In this way, our assessment is not simply of the one particular argument at hand, but rather it is of a whole class of arguments *sufficiently similar* to the one at hand. (Exactly how we explain what we mean by the phrase “sufficiently similar” will determine different formulations of logic.)

Given a symbolic representation of our argument, then, we will want to apply some kind of test. The test should work in such a way that if our original argument was a good one, then the test outputs “good” (and “bad” if the argument was a bad one). The first step in devising such a test will

be to narrow down exactly what it is that we mean by our terms “good” and “bad.”

1.2 Definitions

Let us call an argument **valid** if and only if it is not possible to make its premises true at the same time as its conclusion false. The **premises** are what we’re arguing from, the **conclusion** is what we’re arguing for. (Our definition of validity reveals that we are interested in *deductive* arguments, as opposed to *inductive* ones.) An argument that is not valid, we will call **invalid**. Note that this is meant to coincide with our initial judgments of “good” and “bad.”

Let us call an argument **sound** if and only if (i) the argument is valid, and (ii) its premises are (in fact) true. The collection of all sound arguments, then, is a subset of the collection of all valid arguments. That is, all sound arguments are valid, while not all valid arguments are sound. Note that in our definition of validity, we did not ask whether or not the premises were *actually* true. (If an argument’s premises *were* true and the conclusion could still be false, then it is invalid.)

1.3 Exercises

1. Is each of the following arguments valid or invalid?

(1.1) *All students are mortal. Bertrand is a student. Therefore, Bertrand is mortal.*

(1.2) *Snow is black. The sky is white. Therefore, snow is black and the sky is white.*

2. Is each of the following arguments sound or unsound?

(2.1) *Some chalk is white. Some pens are black. Therefore, some chalk is white and some pens are red.*

(2.2) *All men love sports and cars. Your instructor is a man. Therefore, your instructor loves sports and cars.*

(2.3) *Snow is white. The sky is blue. Therefore, the snow is white and grass is green.*

3. Is it possible for a sound argument to have a false conclusion? Explain.

2 Connectives and truth tables.

2.1 ‘Not’

The first of our logical connectives is ‘Not’. (A **connective** is something that attaches to a sentence—or sentences—making one bigger sentence.) Note that this is our only *one-place* connective, i.e. it “attaches” to a single sentence. Think of the clumsy English expression “It is not the case that.” If I assert: “It is raining,” then my assertion is true if and only if it is actually raining. If I assert: “It is not the case that it is raining,” or “It isn’t raining,” then these are true just in case it is not true that it *is* raining. More abstractly put:

p	¬p
T	F
F	T

This is to say the following. If p is true, then its negation, $\neg p$, is false. Similarly, if $\neg p$ is true, then just p is false. Check that this is correct by trying a couple of example sentences on your own.

Note that many times, in English, ‘not’ sentences are somewhat disguised as contractions. So “Dogs aren’t dumb” says the same thing as the more explicit “It is not the case that dogs are dumb.”

2.2 ‘And’

The rest of our connectives are *two-place* connectives. This is to say that they “connect” two sentences. ‘And’ works in the following way. It is true if and only if each of the parts (called *conjuncts*) are true. Otherwise, the ‘and sentence’ (called a *conjunction*) is false. In table form:

p	q	$p \wedge q$
T	T	T
T	F	F
F	T	F
F	F	F

Check that the values represented in this table make sense along with your understanding of ‘and’ in English.

2.3 ‘Or’

An ‘or’ sentence (called a *disjunction*) is true exactly when any of the parts is true. (We call the parts *disjuncts*.) Note that this definition implies that when both parts of a disjunction are true, so is the disjunction as a

whole. (Our ‘or’ is *inclusive* of this true-true case.) Represented in table form:

p	q	$p \vee q$
T	T	T
T	F	T
F	T	T
F	F	F

Again, compare the values represented above with your “intuitive” understanding of ‘or’ in English. (Sometimes, in English, when we assert a disjunction, we mean to assert that *only one* of the disjuncts can be true. Our wedge does *not* correspond to such usage. We will return to this point later.)

2.4 ‘Only if’

The ‘only if’ corresponds to a conditional assertion in English. Often, this is seen as a sentence of the form “If... then...” (The following forms are equivalent, for our purposes: ‘ p only if q ’ and ‘if p , then q ’.) The key to understanding the conditional is that it is false exactly when the first part (called the *antecedent*) is true and the second part (called the *consequent*) is false. In all other cases, the conditional is true. As a table:

p	q	$p \rightarrow q$
T	T	T
T	F	F
F	T	T
F	F	T

Think of examples, such as the following: “If you do your homework, then you may watch television.” Your parent would have been lying to you—i.e. asserting something *false*—if you *did* do your homework and then you were *not* allowed to watch television. (That is to say, the assertion is false if the antecedent is true and the consequent is false.) The conditional is probably the trickiest of the connectives. One might think it odd, for example, that a conditional with a false antecedent and a false consequent is true. This means that the statement: “If $2+2=5$, then your instructor is the ruler of the world” is *true*. These kinds of examples notwithstanding, there are several good reasons why we take the conditional as represented above. (This way of understanding the conditional is called the “material conditional.”)

2.5 ‘If and only if’

We can think of this last connective basically as a kind of “equal-sign” for truth values. An ‘if and only if’ sentence is called a *biconditional*. As such, when each of the parts (called *limbs*) is true, the biconditional is true. Also, when each of the parts is false, the biconditional is true. When the truth values of the parts differ, then the biconditional is false. Represented in tabular form:

p	q	$p \leftrightarrow q$
T	T	T
T	F	F
F	T	F
F	F	T

One other note about ‘if and only if’. This phrase will often be abbreviated as ‘iff’. You will see ‘iff’ used frequently in definitions. For example: x is a square *iff* x is a rectangle and its sides are all of equal length. See above, throughout, for more examples of this usage of ‘if and only if’.

2.6 Definitions

Let us say that a sentence is **complex** iff contains at least one connective. A sentence that is not complex is called **simple**. A complex sentence (since it is well-formed) will then also contain at least one simple sentence.

The above tables are called **truth tables**. Truth tables lay out all of the possible truth assignments to the component parts of a sentence. For n component parts, there will be 2^n possible truth assignments. (The base here is 2 because there are two truth values: true and false.) So, for a complex sentence built up out of 2 simple sentences, there will be 4 possible truth-assignments. (See, for example, the table for ‘if and only if’ above.) For a complex sentence built up out of 3 simple sentences, there will be 8 possible truth-assignments. Check that each of the tables above captures all of the possible truth value combinations.

2.7 Exercise

1. Suppose that you are told to give a truth table for a new two-place connective. You are given that the connective is true if and only if both parts are false. Let us symbolize the connective with ‘ \downarrow ’. Produce the table for $p \downarrow q$.

3 Regimentation of arguments.

3.1 The strategy.

Our goal is to arrive at a symbolized version of a given argument. Many sentences in arguments are not simple, therefore, some work is sometimes required in order to come up with a symbolized sentence that corresponds to the English version. In particular, one needs to determine what the simple parts of the sentence are, as well as what connectives are used. We will call this process of “coming up with a symbolized version” **regimentation** (or, sometimes, **translation**, though this term can be misleading).

Let us take a single sentence as an example: “The sky is blue and the chalk is white.” What are the simple sentences in this complex one? There are two: we have “The sky is blue” and “The chalk is white.” What is the connective? We have just one: “and”. So, the sentence is a conjunction of two simple sentences. Let us symbolize simple sentences with capital letters. We might assign ‘S’ to “the sky is blue” and ‘C’ to “the chalk is white.” Our completely symbolized sentence is then: $S \wedge C$.

Our strategy for symbolizing entire arguments will be just the same. Each premise will get symbolized, as will the conclusion. (It is important that the same simple sentences get symbolized with the same capital letter, no matter where they appear in the argument.) After doing this, we will be able to use our truth tables to test the symbolized argument for validity (see next section).

3.2 Examples

Example 1: “Dogs bark only if cats sleep. Dogs bark and ticks bite. Therefore, cats sleep.” Take each premise and the conclusion separately. The first premise has two simple components: “Dogs bark” and “Cats sleep.” The connective is the ‘only if’. So symbolizing “dogs bark” as ‘D’ and “cats sleep” as ‘C’, we get the symbolized sentence: $D \rightarrow C$.

The second premise has a new simple sentence as a component: “ticks bite”. Let us symbolize this with ‘T’. The whole second premise will then be symbolized as: $D \wedge T$.

The conclusion has no new components. We shall symbolize it, then, simply as C .

The whole argument, then: $D \rightarrow C, D \wedge T, \text{ therefore } C$.

Example 2: “You will pass this course or you will not be happy. You will be happy. Therefore, you will pass this course.” The first premise has

two simple components. They are “you will pass this course” (‘P’), and the second is “you will be happy” (‘H’). Note that the “not” is *not* part of the simple component: its presence in a sentence means that the sentence is complex. So the first premise is symbolized: $P \vee \neg H$.

The second premise has no new simple components. It will just be symbolized as follows: H .

The conclusion has no new parts. It will be symbolized as follows: P .

The whole argument symbolized, then: $P \vee \neg H, H, \text{ therefore } P$.

3.3 Exercises

Symbolize the following arguments.

(1) *Snow is white. The sky is blue. Therefore, the snow is white and grass is green.*

(2) *The Cubs will win or the White Sox will win. The Cubs will not win. Therefore, the White Sox will win.*

(3) *We will visit Mars only if it is inhabited. It is inhabited. Therefore, we will visit it.*

(4) *If Mars is inhabited, then we will visit it. It is inhabited. Therefore, we will visit Mars or we will visit Jupiter.*

4 Testing an argument for validity.

4.1 Computing Values of Complex Expressions.

First recall our definition of validity. An argument is valid if and only if there is no possibility of making the conclusion false while holding the premises to be true. What we need, then, is a way of surveying the different possible **truth-value assignments**. We already have such a method: truth tables.

What we need, however, is to be able to examine the truth-value assignments of more complex sentences. (All of our previous tables were, we might say, the most basic tables for each of the connectives.) Again, we already have the tools before us. Since all of our complex sentences are built up out of our five connectives (together with simple sentence letters), we can use the basic truth tables to compute the values of the more complex sentences. Let us look at an example.

Example 1. $(A \wedge B) \rightarrow B$.

Set up a truth table for this expression. As it contains two simple sentence letters ('A' and 'B'), we'll need $2^2 = 4$ rows.

A	B	$(A \wedge B) \rightarrow B$
T	T	
T	F	
F	T	
F	F	

The real question here is how to fill in the rest of the table. Notice that the complex sentence is an 'only if' (conditional) statement. It happens to have a complex antecedent, and this antecedent is itself a conjunction. We know how to fill in the table for $A \wedge B$; that's just our basic 'and' truth table. So let's put in that part. (We'll fill it in under the ' $A \wedge B$ '.)

A	B	$(A \wedge B) \rightarrow B$
T	T	T
T	F	F
F	T	F
F	F	F

So now we have the values for the antecedent filled in. We know the values for the consequent; they're simply the column for B from the front of the table. Let's fill that in, under the 'B'.

A	B	$(A \wedge B) \rightarrow B$
T	T	T T
T	F	F F
F	T	F T
F	F	F F

We have just one more step to go, and that's the over-all value for the whole complex expression. What we do is to use the column for the antecedent and the column for the consequent, and we enter the value for the 'only if' using those values together with our basic table for the arrow. That will give us the complete table:

A	B	$(A \wedge B) \rightarrow B$
T	T	T T T
T	F	F T F
F	T	F T T
F	F	F T F

We have all 'T's going down the final column, because there is no instance (row) where the antecedent is true and the consequent is false. (Notice that this means that the sentence is *never* false, no matter what truth values we assign to A and B .)

Example 2. $[(A \leftrightarrow B) \vee \neg A] \wedge (A \vee B)$

A	B	$[(A \leftrightarrow B) \vee \neg A] \wedge (A \vee B)$				
T	T	T	T	F	T	T
T	F	F	F	F	F	T
F	T	F	T	T	T	T
F	F	T	T	T	F	F

This formula is a conjunction. (Another way of saying this: the **main connective** of the formula is the conjunction represented by the wedge. The main connective is above the *last* column you compute.) Hence, the column in bold (under the wedge) represents the table for the expression as a whole. The first column was computed first, then the third. Then the second column could be computed, and this represents the values for the left conjunct. The right conjunct is simply a basic ‘or’ statement. Now we used the left conjunct’s values together with the right’s, and we mark a ‘T’ where both values are true. (Otherwise, we mark an ‘F’.)

4.2 Exercises

Complete truth tables for the following complex expressions.

- (1) $(A \vee B) \rightarrow (A \wedge B)$
- (2) $(\neg A \leftrightarrow B) \vee (\neg B \rightarrow A)$

4.3 Testing *Arguments* for Validity.

Our goal was to test complete arguments for validity. We now have a method for doing this: we list all of the premises, together with the conclusion, at the top of the table. We then compute their values. When this is complete, we examine the table. Are there any rows in which the premises are each true, while the conclusion is false? If so, this row represents a **counterexample**, and the argument is **invalid**. If there aren’t any such rows, then there are no counterexamples, and the argument is **valid**.

Example 3. $(A \vee B), \neg B, \text{ therefore } A$

A	B	$(A \vee B)$	$\neg B$	A
T	T	T	F	T
T	F	T	T	T
F	T	T	F	F
F	F	F	T	F

So this argument is *valid*: there is no row in which the premises come out true and the conclusion comes out false. Let us look at a slightly different example.

Example 4. $(A \vee B), B$, therefore A

A	B	$(A \vee B)$	B	A
T	T	T	T	T
T	F	T	F	T
F	T	T	T	F ←
F	F	F	F	F

So this argument is *invalid*: the arrow at row three indicates the counterexample. In this row, both of the premises are true and the conclusion is false. In other words, we have a counterexample when A is false and B is true.

4.4 Exercises.

(1)–(4): Test the arguments of §3.3, above, for validity.

5 Two common valid argument forms.

5.1 “Modus ponens”

Modus ponens (latin for “affirming mode”) is the name for the following pattern of inference:

$$p \rightarrow q, p, \text{ therefore } q$$

An example: “If you like cake, then you like frosting. You like cake. Therefore, you like frosting.”

5.2 “Modus tollens”

Modus tollens (latin for “denying mode”) is the name for the following pattern of inference:

$$p \rightarrow q, \neg q, \text{ therefore } \neg p$$

An example: “If you like cake, then you like frosting. You don’t like frosting. Therefore, you don’t like cake.”

6 Two fallacious argument forms.

6.1 “Affirming the consequent”

The following argument form is *not* a valid pattern of inference:

$$p \rightarrow q, q, \text{ therefore } p$$

An example: “If you like cake, then you like frosting. You like frosting. Therefore, you like cake.”

6.2 “Denying the antecedent”

The following argument form is also *not* a valid pattern of inference:

$$p \rightarrow q, \neg p, \text{ therefore } \neg q$$

An example: “If you like cake, then you like frosting. You don’t like cake. Therefore, you don’t like frosting.”

6.3 Exercises.

(1) Show that *modus ponens* and *modus tollens* are each indeed valid patterns of inference.

(2) Show that “affirming the consequent” and “denying the antecedent” are indeed *not* valid patterns of inference. (Do this by giving an assignment of truth-values to the letters p and q that make the premises of the argument form true while making the conclusion false.)

7 Use and Mention

7.1 Motivation.

There is a difference, in our language, between *using* an expression and *mentioning* one. It should be obvious that it is important, for whatever purposes we have, that we know what we're saying with our words. When I mention an expression, I am talking about the expression itself—as opposed to talking about whatever it is to which the expression refers. Here is an example of mentioning an expression:

The first letter of “dog” is a consonant.

We use the quotation marks to designate the fact that we are mentioning the expression, instead of using it. (We can think of the quotation marks as providing us a kind of name for the word, enabling us to talk about it.) Consider the same sentence without the quotation marks around “dog”:

The first letter of dog is a consonant.

This sentence is meaningless! It purports to talk about something called “the first letter of dog”. This is why it is important to distinguish between use and mention: so that our sentences at least make sense. We can evaluate the philosophical adequacy of what we're talking about later—but if we're not even making sense with our words (if aren't clear about what it is that we're talking about) then our philosophical conversation can't even really begin.

Notice the use of quotation marks in the sentence before this last example sentence. There, I wished to talk about the sentence not having quotation marks around a certain word, so I needed to use quotation marks when referring to that word. (Otherwise, I would have been talking about some actual dog that happened to have quotation marks hovering above *it...*)

We shall use single quotes (‘’) and double quotes (“”) much as we use parentheses () and brackets []. Single quotes will be used to mention expressions inside of an expression that is already mentioning an expression. An example of this:

“The first letter of ‘dog’ is a consonant” is true.

7.2 Exercises.

Place quotation marks in the following sentences in order to have them make sense grammatically and factually.

- (1) Snow is white is true if and only if snow is white.
- (2) Grass is green is true but grass is only five letters long.
- (3) Use to mention to mention to mention.

8 Some concluding exercises.

(1) (Jeffrey, p.20; also Smullyan) Knaves always lie, knights always tell the truth, and in Transylvania, where everybody is one or the other (but you can't tell which by looking), you encounter two people, one of whom says "He's a knight or I'm a knave." What are they? (Hint: you don't necessarily need to try to regiment this argument; rather, try reasoning through the different possibilities.)

(2) Explain the process of regimenting an argument.

(3) Symbolize the following argument and test it for validity:

"If Bob is a snark, then Bob is a dwark. Bob is not a dwark or Bob is a thwark. Bob is not a thwark. Therefore, Bob is not a snark."

(3.1) Is the argument in (3) sound or unsound? Explain your answer.

(4) Place quotation marks in the following, in order to have it make sense grammatically:

I used to use to mention to mention to mention.

9 References

Hart, W.D. "Some exercises on use and mention." (unpublished)

Jeffrey, Richard *Formal Logic: Its Scope and Limits* 3rd Edition (McGraw-Hill: 1991)

Klenk, Virginia *Symbolic Logic* 3rd Edition (Prentice Hall: 1994)

Quine, W. V. O. *Philosophy of Logic* 2nd Edition (Harvard: 1986)

Quine, W. V. O. *Elementary Logic* Revised Edition (Harvard: 1980)