A Brief Introduction to Logic

Paul A. Gregory
Washington and Lee University

1 Belief and Critical Thought

Viewed from a certain angle, philosophy is about what, if anything, we ought to believe. Much philosophy, such as moral theory or social and political philosophy, has to do with how we should act and how we should structure our society's institutions, so one might say that viewing philosophy as all about beliefs ignores these aspects. But how we decide to act or how we hope to change society ultimately comes down to what we believe about the nature of the world, about our nature as human beings, about the facts of a particular situation, and what we value—what we believe is good and right. Thus, for present purposes, it is accurate enough to say that philosophy is all about what we should believe.

There are many sources for our beliefs—the influence of our parents and friends, newspapers, television, internet, government, churches, schools, textbooks, our own observations, and our own thought processes—but not all these sources are equally good, nor is any of them good all the time. Most of us form most of our beliefs without ever consciously attending to the fact that we have acquired a belief. Many of the beliefs attained through casual observation and most of the teaching and socialization of childhood is simply absorbed without our realizing it. Despite our best efforts, this uncritical acceptance of claims and formation of beliefs continues through adulthood. But such uncritical thinking can be a dangerous thing, for it enslaves us to the influence and manipulations of other individuals and institutions. Thus, since our beliefs (the claims we accept) constitute our view of the world and of ourselves, and affect how we act, it is important to examine more carefully the beliefs we hold. This applies both to large scale philosophical issues (Does God exist? Do I have a soul? Am I free? Does it matter?), and to small scale mundane issues (Should I buy this car? Should I believe what she said? How should I dress?).

This is what we will spend the rest of the semester doing—examining various of our fundamental philosophical beliefs about the world, ourselves, and our values. In this brief piece on logic, I want to introduce a framework of technical terms which we will often use when examining claims and the arguments or evidence offered in support of them.

2 Logic: What and Why?

What is Logic?

Logic is the study of argument. In particular, logic is the study of criteria for distinguishing successful from unsuccessful arguments and the study of methods for applying those criteria.

By 'argument' I do not mean a shouting match or angry disagreement (though sometimes these accompany the sort of argument which interests us). Rather, an argument is a set of statements, some of which—the *premises*—are supposed to support, or give reasons for, the remaining statement—the *conclusion*. Such statements, and so the argument they constitute, may be written, spoken, thought, or otherwise communicated or privately considered.

Usually the intent behind an argument is to produce understanding or conviction in oneself or another. However, by 'successful argument', I do *not* mean one which merely succeeds in persuading the reader or listener or thinker. Due to the foibles of human psychology and/or the use of manipulative rhetorical devices, people can be persuaded by arguments even though the premises do not genuinely support the conclusion. Moreover, people can fail to be convinced by even the tightest and clearest reasoning. Strategies for detecting and analyzing manipulative rhetoric and strategies for overcoming natural human obtuseness will not be the focus of this brief introduction—those would require a much more extended treatment. Here we shall focus on the core principles which any such critical strategies presuppose: the criteria for successful argument.

The basic idea behind these criteria is that the premises should genuinely support the conclusion—as opposed, for example, to merely inspiring a new, or reinforcing an existing, emotional commitment to the conclusion. To achieve genuine support, the premises should be related to the conclusion in a way which either guarantees or makes probable the preservation of truth from premises to conclusion. In other words, in a successful argument, if the premises are true, then the conclusion is either guaranteed to be true or likely to be true. These rough statements will require much clarification and elaboration. But, before moving on, it is worth considering why one might want to study logic.

Why Study Logic?

There is a great deal of instrumental value in studying logic—logic is *good for* other desirable things. First, as alluded to above, skill in logical analysis and evaluation is the core of critical thinking. So, increasing your logical abilities will lead to a corresponding increase in your overall critical ability. You will be better able to guard against psychological fallibility and manipulative rhetoric—your own as well as others'.

While a basic understanding of some logic will not determine for you what you ought to believe, it is an essential tool for being a strong critical thinker. And whether you are writing a paper in philosophy, having a discussion with a friend, watching the news, or buying a computer, it is always to the good to be a strong critical thinker.

Of course, beyond instrumental value, there is a great deal of intrinsic value—the study of logic is good *in itself*. First, there is the discovery and learning of new things. Then there is aesthetic and intellectual delight in abstract structures and systems. Less grandiose, perhaps, is the enjoyment of puzzle solving and rising to meet an intellectual challenge. This is not to say that appreciating logic as good in itself is required to do well in philosophy. Far from it. But it will help you to understand some logic.

3 Arguments, Forms, and Truth Values

Now it is time to get down to business. To start out, let's officially define some of the terms used in the previous section, as well as some new terms:

Logic:
: *Logic* is the study of (i) criteria for distinguishing successful from unsuccessful argument, (ii) methods for applying those criteria, and (iii) related properties of statements such as implication, equivalence, logical truth, consistency, etc.

Part (iii) of the definition of 'logic' will not be discussed at length here, but I include it here for the sake of completeness.

Statement:
: A *statement* is a declarative sentence; a sentence which attempts to state a fact—as opposed to a question, a command, an exclamation.

Truth Value:
: The *truth value* of a statement is just its truth or falsehood. At this point we

> make the assumption that every statement is either true (has the truth value true) or false (has the truth value false) but not both. The truth value of a given statement is fixed whether or not we *know* what that truth value is.

For variety, throughout most of the discussion ‘statement’, ‘sentence’, and ‘claim’ will be used interchangeably to refer to declarative sentences. Moreover, our assumption regarding truth values—that all statements are either true or false and not both—is one that may reasonably be questioned. Issues of vagueness, ambiguity, subjectivity, and various sorts of indeterminacy may lead us to think that some statements are neither true nor false, or both true and false, or “somewhere in between”. There are logics which attempt to formalize such intuitions by countenancing three or more values (as opposed to our two), and treating compound claims and arguments in correspondingly more complex ways than we will. To genuinely appreciate such logics, however, it is important to have an understanding of classical two-valued (monotonic, or bi-valent) logic. So we’ll rely on this assumption for our purposes.

Argument:

> An *argument* is a (finite) set of statements, some of which—the *premises*—are supposed to support, or give reasons for, the remaining statement—the *conclusion*.

When we encounter an argument in the course of reading or during discussion the premises and conclusion may come in any order. Consider the following versions of a classic example:

- Socrates is mortal, for all humans are mortal, and Socrates is human
- Given that Socrates is human, Socrates is mortal; since all humans are mortal
- All humans are mortal; Socrates is human; therefore, Socrates is mortal

Three statements are involved in each of the above examples (two premises and a conclusion), and despite the fact that they appear in different order in each one, all three examples express the same argument. For the sake of clarity we will often transcribe arguments into what is called *standard form*—we list the premises, draw a line, then write the conclusion:

$$\begin{array}{l} \text{Premise 1} \\ \text{Premise 2} \\ \quad\vdots \\ \text{Premise } n \\ \hline \text{Conclusion} \end{array} \qquad\qquad \begin{array}{l} \text{All humans are mortal} \\ \text{Socrates is human} \\ \hline \text{Socrates is mortal} \end{array}$$

Despite the variability of statement order when not expressed in standard form, a good writer usually makes clear which sentences are premises and which is the conclusion. This is usually done through contextual clues, including the indicator words/phrases I used above. Below are two brief lists of conclusion and premise indicators:

Premise Indicators:
as, since, for, because, given that, for the reason that, inasmuch as

Conclusion Indicators:
therefore, hence, thus, so, we may infer, consequently, it follows that

So it is usually a relatively simple task to put arguments into standard form. Of course, if the argument is long and complex, with sub-conclusions acting as premises for further conclusions, analysis can get messy.

In addition to standard form, the notions of argument *form* and *instance* will be important in what follows. An *argument form* (sometimes called a schema) is the framework of an argument which results when certain portions of the component sentences are replaced by blanks or schematic letters. An *argument instance* is what results when the blanks in a form are appropriately filled in. For example, the argument presented above is an instance of the following form:

All F are G
x is F
---
x is G

Here 'F' and 'G' are placeholders for predicate phrases, while 'x' is a placeholder for a name. The form was generated by replacing 'humans' with 'F', 'mortal' with 'G', and 'Socrates' with 'x'. So we have here our first use of symbols as a means to capturing logical form. I trust it is fairly intuitive at this point. With little or no effort you will see that the following is also an instance of the above form:

All pop-stars are attention-starved
Mr. Green Jeans is a pop star

Mr. Green Jeans is attention-starved

Here 'pop-star' goes in for 'F', 'attention-starved' for 'G', and 'Mr. Green Jeans' for 'x'. The result is another argument of the same form.

Argument Form and Instance:
An *argument form* (or schema) is the framework of an argument which results when certain portions of the component statements are replaced by blanks, schematic letters, or other symbols. An *argument instance* is what results when the blanks in a form are appropriately filled in.

Now I will divide criteria for evaluating arguments into two basic types. We will have similar but different things to say about what counts as "success" for each type. As mentioned previously, the two types are Deductive and Inductive. These are discussed in sections 4 and 5, respectively.

4 Deductive Criteria

Deductive criteria, roughly speaking, require that the premises guarantee the truth of the conclusion. There are two questions we want to ask when applying deductive criteria. One, the question of *validity*, has to do with the connection between the premises and conclusion. The other, the question of *soundness*, has to do with the truth values of the premises. First, validity:

Deductive Validity, Invalidity:

> An argument (form) is *deductively valid* if and only if it is NOT possible for ALL the premises to be true AND the conclusion false. An argument (form) is *deductively invalid* if and only if it is not valid.

I.e., an argument is valid if and only if the assumed truth of the premises would guarantee the truth of the conclusion—if you mull it over you'll see that these say basically the same thing; though the statement above is our official definition of validity. (Where no confusion threatens, I will drop the 'deductive' and 'deductively' from my talk of validity.) Note also that the definition of validity applies both to individual arguments and to argument forms. This will be made clear in what follows.

In typical contexts we want more than just validity from our arguments. For, even if they are valid, this means nothing about the truth of the conclusion, unless the premises are true as well—i.e., unless the argument is *sound*.

Soundness:

> An argument is *sound* if and only if it is deductively valid AND all its premises are true.

To begin understanding and distinguishing these two criteria, consider the group of arguments in Table 1. Upon reflection, you should be able to see that all three arguments of **Form 1** are valid, and that *this has nothing to do with the actual truth values of the component claims*. That is, despite the actual truth value of the component statements in each instance, there is no way (it is NOT possible) that the premises could ALL be true AND the conclusion false. Take **1D**, for example. Despite the fact that all the claims are false, it has to be the case that IF the premises

	All premises True	At least one premise False
Conclusion True	**1A** All whales are mammals All mammals are air-breathers All whales are air-breathers Valid and Sound	**1B** All whales are fish All fish are air-breathers All whales are air-breathers Valid but Unsound
Conclusion False	**1C** Not Possible!	**1D** All whales are reptiles All reptiles are birds All whales are birds Valid but Unsound

Form 1
All F are G All G are H All F are H Valid Form

Table 1: Argument Form 1 with Instances

were true, then the conclusion would be true as well—just imagine the premises are true; could the conclusion be false? NO, so it's valid.

As you will have noticed, **1A**, **1B**, and **1D** are all of the same form. Indeed, it is in virtue of having this form that each of the three instances are valid. Any argument in which we consistently substitute predicate phrases for the place-holders F, G, and H will be a valid argument. Thus, **Form 1** is a valid argument form and **1A**, **1B**, and **1D** are valid arguments because they are instances of a valid form. Validity is a question of form. That is why the actual truth values of the component claims are irrelevant. What is relevant is whether the form is such that it is not possible to have all premises true and the conclusion false.

Another way of putting this is that valid arguments are truth preserving. A good metaphor for this is plumbing: if you hook up the pipes correctly (if your argument has a valid form) you know that if you put water in at the top (true premises), you'll get water out at the bottom (true conclusion). But it doesn't actually matter whether you do, indeed, put any water in—the pipes are hooked up correctly (the argument is valid) whether or not you put any water through them (whether or not the premises are true).

Note that of the three arguments of **Form 1**, only **1A** is also sound. This is because, in addition to being valid, it has premises which are actually true.

What, then, of **1C**? There can be no such instance of **Form 1**. It is impossible, the argument form is valid. Compare the arguments of **Form 2**, **2C** in particular (in Table 2).

	All premises True	At least one premise False
Conclusion True	**2A** Some animals are frogs Some animals are tree-climbers Some frogs are tree-climbers Invalid	**2B** Some fish are frogs Some fish are tree-climbers Some frogs are tree-climbers Invalid
Conclusion False	**2C** Some animals are frogs Some animals are birds Some frogs are birds Invalid	**2D** Some fish are frogs Some fish are birds Some frogs are birds Invalid

Form 2

Some *F* are *G*
Some *F* are *H*
Some *G* are *H*

Invalid Form

Table 2: Argument Form 2 with Instances

Now look at the four arguments of **Form 2** in Table 2. First, note that they are all invalid. That is, despite the actual truth values of the component sentences, in each instance it is possible for all the premises to be true and the conclusion false. **2A**, **2B**, and **2D**, will require some imagination, but you will see that you can consistently imagine all the premises true and the conclusion false in each case.

2C takes no imagination at all. Here we have an instance of the form in which it is obvious that the premises are all actually true and the conclusion is actually false. Obviously this particular argument instance is invalid. Since the premises are all true and the conclusion is false, it is *possible* for all the premises to be true and the conclusion false (actuality implies possibility). Moreover, since validity is a matter of form, once we have an instance like **2C** before our eyes, we know that the form is invalid, and so is any instance (such as **2A**, **2B**, or **2D**). Notice that it was precisely the C-type instance (all premises True, conclusion False) that the valid **Form 1** lacked. When we find such a C-type argument we are said to have found a *counterexample* to the argument form. Valid arguments do not have counterexamples.

Counterexample:

> A *counterexample* to an argument (form) is an argument instance of exactly the same form having all true premises and a false conclusion. Production of a counterexample shows that the argument form and all instances thereof are invalid. (Failure to produce a counterexample shows nothing, however).

Stretching the plumbing metaphor a bit, the counterexample, **2C**, illustrates that if your pipes are not correctly connected (the argument form is invalid) there is no guarantee that putting water in at the top (true premises) will result in water coming out at the bottom (true conclusion). You might get lucky and have water come out at the bottom (you might end up with a true conclusion as in **2A**), but , as shown by **2C**, you might not. And you should never take chances with your plumbing.

It is important to note, however, that the inability to produce a counterexample does not show validity. It may just be that we are (perhaps momentarily) not clever enough to fill in the form so as to have true premises and false conclusion. For determining validity and invalidity, it would be nice to have some method more systematic and reliable than our unschooled imaginations—the formal symbolic approach of subsequent chapters will give us a number of these.

Lastly, because an argument must be valid to be sound, and none of the instances of **Form 2** is valid, none of them is sound—in fact, once we determine that an argument is invalid, we don't bother with the question of soundness. Figure 1 presents a flowchart for assessing validity and soundness.

Here are some points to remember about deductive validity.

- Validity is a question of truth preservation, and this is a question of form, so the actual truth values of the premises and conclusion are irrelevant.
- *All true premises and true conclusion,* ***do not*** *make a valid argument!* See **1B**, **1D**, and (especially) **2A**.
- If an argument is valid and all its premises are true, then it is sound.
- Soundness *does* have to do with the actual truth value of the premises. Thus, it is not a matter of form.
- We can see that a particular argument, and all arguments of the same form, are invalid either by consistently imagining that all the premises are true and the conclusion false, or by finding a counterexample (an instance which actually does have all true premises and a false conclusion).

The following are some *valid* forms with example instances. (Here '*P*' and '*Q*' are place-holders for sentences.)

Disjunctive Syllogism:

Either *P* or *Q*	Either Pat is a man or Pat is a woman
not-*Q*	Pat is not a woman
P	Pat is a man

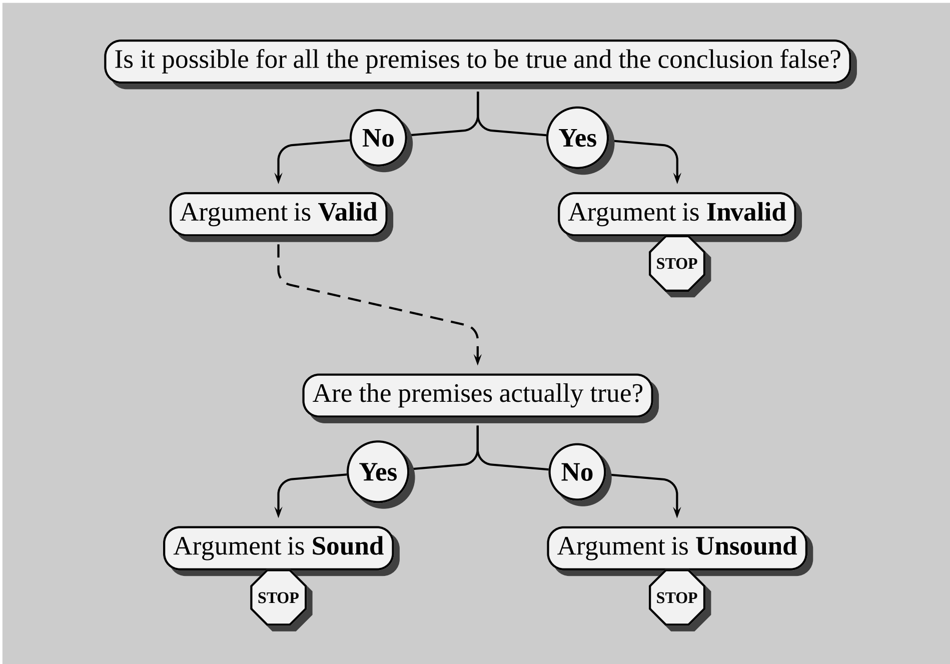

Figure 1: Evaluating Deductive Arguments

Reductio Ad Absurdum:	
Assume *P*	
deduce a contradiction:	Suppose Pat is a mother
	All mothers are women
Q	So, Pat is a woman
not-*Q*	But, Pat is a man, not a woman
not-*P*	So, Pat is not a mother
Modus Ponens:	
If *P*, then *Q*	If Pat is a mother, then Pat is a woman
P	Pat is a mother
Q	So, Pat is a woman
Modus Tollens:	
If *P*, then *Q*	If Pat is a mother, then Pat is a woman
not-*Q*	Pat is not a woman
not-*P*	So, Pat is not a mother

The following two, though tempting, are *invalid*.

Denying the Antecedent *(invalid)***:**

If P, then Q	If Pat is a mother, then Pat is a woman
not-P	Pat is a not a mother
not-Q	So, Pat is not a woman

Affirming the Consequent *(invalid)***:**

If P, then Q	If Pat is a mother, then Pat is a woman
Q	Pat is a woman
P	So, Pat is a mother

4.1 Other Deductive Properties

There are other deductively relevant properties of individual statements and sets of statements which can be analyzed as a matter of form. Thus the earlier statement, "Logic is the study of argument", is a useful oversimplification. Logic is also concerned with such things as logical truth, equivalence, consistency, and entailment, to name just a few.

Consider the following three statements:

(1) Xeno had twelve toes
(2) Either Xeno had twelve toes or Xeno did not have twelve toes
(3) Xeno had twelve toes and Xeno did not have twelve toes

Unless you have some very unique historical insight, you probably do not know whether (1) is true or false. But what about (2)? Whether or not you even know who Xeno is, you should be able to tell that (2) is true. Indeed, you might be thinking to yourself, 'It has to be true; it just can't be false'. And you would be absolutely correct. (2) is *logically true*. It's logical structure is such that it is not possible for it to be false, it must be true. For, either the guy had twelve toes, in which case (2) is true; or he didn't, but in this case (2) is true again. So in any case the statement is true.

With (2), as with many logical truths, we can see the truth of the statement even if we do not know much about the subject matter. This is because their truth is entirely a matter of logical form. To see this, Consider the form of (2), and another instance of the same form (here 'P' is a place holder for statements):

(2′) Either 'P' or not 'P'
(2″) Either Xeno was an avid distance runner or Xeno was not an avid distance runner

You may feel that sentences such as (2) and (2″) don't really say that much. They include all the possibilities without really asserting any one of them in particular in particular. You would be in good company—a philosophically important theory of logical truth claims that all logical truth are like this. (2) and (2″) are instances of what are sometimes called tautologies. The word has both the technical use when applied to logical truths, and the original, more common, use when applied to an obvious or redundant statement. But this is just the nature of some logical truths. They may not be very informative or useful for communication, but they are, logically speaking, very interesting statements. Moreover, not all logical truths are as obvious and easy to spot as (2) and (2″); logical truths can be quite complicated and difficult to recognize. But we will develop formal symbolic methods to test for them.

Return now to (3). Here, again, you can tell the truth value without really knowing anything about Xeno. Briefly, because it says that he both did and did not have twelve toes, it must be false. The logical structure of the statement forces it to be false. (3), perhaps unsurprisingly, is what we call *logically false*. Again, it is a matter of form. See if you can give the form of (3) and another instance of the same form (in the style of (2′) and (2″)).

Finally, reconsider (1). Very likely it is false—I have no idea—but it certainly could be true. That is, there is nothing about its *logical* form which either requires that it be true or requires that it be false. It could be either. In this case we call the statement logically contingent. Here are some definitions:

Logically True:
A statement is *logically true* if and only if it is not possible for the statement to be false.

Logically False:
A statement is *logically false* if and only if it is not possible for the statement to be true.

Logically Contingent:
A statement is *logically contingent* if and only if it is neither logically true nor logically false; i.e., it is both possible for the statement to be true, and possible for the statement to be false.

Now consider some pairs of sentences:

(4) In addition to having twelve toes, Xeno was an avid distance runner
Xeno was an avid distance runner, even though he had twelve toes

(5) Xeno finished some of his races
Xeno did not finish any of his races

(6) Xeno had twelve toes
Xeno never finished a race

The statements in pair (4) do not say *exactly* the same thing. There is a clear difference in rhetorical tone and emphasis, especially with regard to his polydactyly[1]. However, we will not be concerned with such things a tone, emphasis, connotation. From a logical point of view both statements assert that Xeno had twelve toes and that Xeno was an avid distance runner. Thus, we consider the two sentences *logically equivalent*, meaning that their forms are such that they will always have the same truth value.

Contrast the equivalence of pair (4) with the situation in pair (5). Here, it seems the two statements are saying the exact opposite of one another—the second in the pair contradicts the first (and *vice versa*). We call such a pair *logically contradictory*. Their logical forms are such that they always have opposite truth values.

Pair (6), in contrast to pairs (4) and (5), has nothing very special about it. Unlike pair (4), the statements of (6) could differ in truth value. Unlike pair (5), they do not have to differ in truth vale. Hence the pair in (6) is neither logically equivalent, nor logically contradictory.

Logically Equivalent:
A pair of statements is *logically equivalent* if and only if it is not possible for the statements to have different truth values.

Logically Contradictory:
A pair of statements is *logically contradictory* if and only if it is not possible for the statements to have the same truth values.

In addition to looking at individual statements and pairs of statements, we can focus on sets of statements. A set, roughly, is a grouping of any number of objects, considered as a group. A set can contain zero, one, two, three, four,..., an infinite number of members.[2] Below are a couple of sets of statements. Notice that we list the members of the sets (the statements, in this case), separated by commas, with the list enclosed in curly braces.

(7) { Everyone with twelve toes requires custom racing sandals, Xeno had twelve toes, Xeno did not require custom racing sandals }

[1] Dodecadactyly!?

[2] The set with zero objects is called, aptly enough, the empty set (or null set). For more on sets and set theory, see **Chapter ??**.

(8) { Xeno did not require custom racing sandals, Xeno was an avid distance runner, Xeno had twelve toes }

If you think carefully about set (7), you may notice something is not quite right. If the first and second members of the set are true, then Xeno must have required custom racing sandals. But the third member of the set denies this. If the first and third are true, then Xeno must not have had twelve toes; but the second statement claims that he does. Finally, if the second and third are true, then the first must be false. In brief, there is no way for all three of the members of set (7) to be true. The set is *logically inconsistent*. No matter what, at least one of the statements must be false. Possibly more than one is false—we don't really know the truth values. But even without knowing that, we can see that at least one must be false.

Set (8) is a different story. Again, we don't know the truth values of the statements (I suspect only the first is true), but it is clear that they could all be true. That is, it is possible that every member of the set is true (collectively, not individually). Hence, the set is *logically consistent*.

Logically Consistent:

A set of statements is *logically consistent* if and only if it is possible for all the statements to be true.

Logically Inconsistent:

A set of statements is *logically inconsistent* if and only if it is not possible for all the statements to be true.

Note that these are complimentary properties—a set of statements is logically inconsistent if and only if it is not logically inconsistent.

The final property we shall discuss here is *entailment*. I'll start with the definition:

Logically Entails, Logically Follows:

A set of statements *logically entails* a target statement if and only if it is NOT possible for every member of the set to be true AND the target statement false. We also say that the target statement *logically follows* from the set.

First, note that this is very close to the definition of deductive validity. Indeed the two concepts are almost identical, with just a few theoretical differences (see below). Here is an example:

(9) { Everyone with twelve toes requires custom racing sandals, Xeno had twelve toes } *entails* Xeno required custom racing sandals

(10) { Xeno was an avid distance runner, Xeno had twelve toes } *does not entail* Xeno never finished a race

The claim of entailment in (9) is correct. There is no way for both members of the set to be true and the target statement false. So the target 'Xeno required custom racing sandals' follows from (or is entailed by) the set. The denial of entailment in (10) is likewise correct. It is clearly possible to be a twelve-toed avid runner who finishes at least one race. I.e., it *is* possible for all members of the set to be true and the target false. So the set does not entail the target (the target does not follow from the set).

Again, this notion is very close to that of deductive validity of an argument. If we take the premises of an argument as the set of statements and the conclusion as the target, then we can say that an argument is deductively valid if and only if the premises logically entail the conclusion. There are some differences, however. First, there is the obvious difference in the way in which the question or claim is posed. With validity, we have premises and a conclusion; with entailment, a set and a target. This would be merely cosmetic, if it were not for the second difference. Second, arguments by definition have a finite number of premises, whereas a set may, in principle, contain an infinite number of statements. This could, for all we know initially, make an important difference. Unexpected things often occur when dealing with infinities. Perhaps there are statements entailed by infinite sets, but which do not follow from any finite set of premises (and so they cannot be the conclusion of any valid argument). As it turns out, a metalogical result called the Compactness Theorem shows that this is not the case. The two concepts do coincide. This is a minor, though significant, result of metalogic. Moreover, the use of the 'entails' and 'follows from' terminology is just as common as the use of 'validity', so it is important for you to be familiar with it.

Note that the definitions of this section make frequent appeal to possibility, necessity, and impossibility—just as the definition of validity did. Again, I have relied on fairly simple examples and the reader's intuitive sense of possibility and logical correctness. Formal symbolic logic (which we will not be pursuing here) studies systematic ways to state and test for these properties. The resources of artificial symbolic languages enable the analysis of form, clear definition of the properties of interest, and the development of methods for exploring those properties. Such would be the focus of a course in formal symbolic logic.

5 Inductive Criteria

Traditionally, deductive reasoning was said to be that which proceeds from the general to the particular, while inductive reasoning was said to move from the particular to the general. But this is incorrect. Some deductively successful arguments have particular premises and particular conclusions. Moreover, some inductively successful arguments invoke general premises and/or arrive at particular conclusions. I, therefore, dispense with this traditional way of making the distinction.

I will distinguish inductive from deductive criteria in terms of the sort of support the premises are required to give the conclusion.

Again, there are two questions we want to ask when applying inductive criteria. One, the question of *strength*, has to do with the connection between the premises and conclusion. The other, the question of *cogency*, has to do with the truth values of the premises. First, strength:

Inductive Strength:

> An argument is *inductively strong* to the degree to which the premises provide evidence to make the truth of the conclusion plausible or probable. If an argument is not strong, it is *weak*.

Note the contrast with deductive validity, which requires that premises guarantee the truth of the conclusion. Here, inductive strength is a matter of the degree of plausibility or probability. Also in contrast to the definition of validity, the definition of strength[3] does not apply to argument forms, but only to individual instances. As we shall soon see, this is because strength is not at all a matter of form.

Cogency:

> An argument is *cogent* if and only if it is inductively strong AND all the premises are true.

With respect to strength, cogency plays a role analogous to that which soundness plays with respect to validity. In both cases we have a question about the connection between premises and conclusion (validity or strength), and then a question about the truth value of the premises.

Consider the following examples:

> This bag has 100 marbles in it
> 80 of them are black
> 20 of them are white
> ______________________________
> The next marble I pick will be black

[3] Where confusion does not threaten, I will often omit 'inductive' and 'inductively' when speaking of strength.

It is 5pm on Monday
But the mail has not come yet
The mail carrier is almost never late
It must be a holiday

In neither of these cases do the premises guarantee the truth of the conclusion. So how successful are these arguments? Well, for the first one we have a pretty good idea: it is quite strong. Barring unforeseen happenings, we would rate the probability of the conclusion at 80%.[4] For the second one, however, it is unclear. It seems pretty strong, but that assessment is vague. And this often is the case with non-statistical assessments of strength. Except where the argument is clearly very weak, often we can only give a vague assessment of its strength.

This is a point to remember about deductive versus inductive criteria: deductive validity is like an on/off switch, an argument is either valid or invalid (and not both); but inductive strength is a matter of degree. Moreover, unlike validity, the strength of an argument is not simply a matter of form. Though form is often relevant to assessing the inductive strength of an argument, we will see below that it is never decisive. With inductive strength, more than form needs to be taken into account.

Consider these two common and simple forms of induction:

Induction by Enumeration:

A_1 is F
A_2 is F
⋮
A_n is F
All As (or the next A) will be F

57 trout from Jacob's Creek were all infected with the RGH virus
So, all trout (or the next trout found) in Jacob's Creek will be infected

Argument by Analogy:

A is F, G, H
B is F, G, H, and I
A is I

My car is a 1999 Toyota Camry
Sue's car is a 1999 Toyota Camry and it gets over 30 miles per gallon
So, my car will get over 30 mpg

With enumeration, generally speaking, the larger the sample, the stronger the argument. As the number of observed examples exhibiting the target property, F,

[4] Of course, if we change the conclusion to 'There is an 80% chance the next marble I pick will be black', then we have a deductive statistical argument.

increases, so does the likelihood of the conclusion (unless the population is infinite). Moreover, the narrower (or more conservative) the conclusion, the stronger the argument. For example, it is a narrower conclusion, and so a safer bet, that the next fish will be infected, than that all fish are (perhaps there are a very small number of resistant fish). With arguments by analogy, strength tends to vary with the number and relevance of shared properties (F, G, and H, three is not a required number). The more the two objects (or groups) have in common, and the more relevant those properties are to the target property, I, the more likely the object in question will also have the target property. In assessing the strength of analogical arguments it is also important to consider relevant dissimilarities. If present, relevant dissimilarities weaken the argument.

But inductive arguments are not so simple. Here I will just illustrate a few difficulties. Consider the following, which are of exactly the same forms as the above arguments:

> The 13,000 days since my birth have all been days on
> which I did not die
> ---
> So, all days (or the next day) will be a day on which
> I do not die

> I like peanuts, am bigger than a breadbox, and have two ears
> Bingo the elephant likes peanuts, is bigger than a breadbox,
> has two ears, and has a trunk
> ---
> So, I have a trunk

Note that neither of these arguments is particularly strong. Despite the rather large sample size in the induction, the conclusion that I will live forever is extremely unlikely, and while the conclusion that I will live through the next day seems to be stronger (because it is narrower), it is not particularly comforting... The analogical argument involving the elephant is obviously ludicrous, mainly because the similarities I have cited are largely irrelevant to the question of my having a trunk. Again analogical arguments are stronger when the similarities cited are relevant to the target property, and when there are few relevant dissimilarities.

Determining the relevance of similarities and dissimilarities (as well as the question of the strength of enumerations) depends to a large degree on background knowledge—knowledge which is often left unstated, but which, when made explicit, or inserted as a new premise, may strengthen or weaken the argument. In some cases (say when hypothesizing about the effect on humans of a drug tested on mice) we may not be entirely sure how relevant the similarities and dissimilarities are, and this affects our assessment of the strength of the argument. This issue of background knowledge is part of why strength is not a question of form; more is involved than just the statements which appear in the argument.

Moreover, the emergence of new evidence can radically alter our assessment of inductive strength. If an argument is deductively valid, this status cannot be changed by the introduction of additional evidence in the form of further premises.[5] When applying inductive criteria, however, the introduction of further evidence in the form of additional premises can increase or decrease the strength of the argument. (Imagine pointing out relevant dissimilarities between me and Bingo, or, in the argument about the mail carrier, adding the evidence that the roads are flooded.)

Here are some points to remember:

- Unlike deductive validity, inductive strength is a matter of degree, not an all-or-nothing, on/off switch.
- Unlike deductive validity, inductive strength is *not* a matter of form.
- Background knowledge and additional information is relevant to the assessment of strength.

For enumerations and analogies:

- The larger the sample size or comparison base group, the stronger the argument.
- The narrower or more conservative the conclusion, the stronger the argument.
- The greater the number of (relevant) similarities, the stronger the argument.
- The fewer the number of (relevant) dissimilarities, the stronger the argument.

Abduction

Abduction is a category of reasoning subject to inductive criteria. It deserves special mention because of its ubiquity in daily and scientific reasoning, and because it makes special appeal to the notion of explanation. Criteria for abductive success are essentially inductive criteria. Thus, successful abduction carries no guarantee of truth preservation, strength is a matter of degree rather than a matter of form, and it is subject to reassessment in light of new evidence.

[5] Even if we introduce a new premise which contradicts one of the old premises! The reason is that if the premises contradict one another, then it is not possible for them all to be true. So, it is not possible for all premises to be true AND the conclusion false. So, the argument is valid. Indeed, every argument with contradictory or inconsistent premises is valid. Note, however, that in such "degenerate" cases of validity the argument can never be sound.

But there is an important difference between abduction and other sorts of inductive attempts. Typical successful induction arrives at a conclusion which is probable when we assume the truth of the premises. In addition to this, abductive reasoning is explicitly aimed at *explaining* the truth of the premises. Abduction tries to answer the question of *why* something is the way it is. Thus, abductive reasoning is often described as *inference to the best explanation*. As a result, assessing the strength of an attempt at induction will involve assessing whether, and how well, the conclusion explains the premises.

Abduction:

> *Abduction* or *abductive reasoning*, also known as *inference to the best explanation*, is a category of reasoning subject to inductive criteria in which the conclusion is supposed to explain the truth of the premises.

The example involving the mail carrier (page 17) is a miniature abductive argument. The conclusion that it is a holiday, would, if true, explain the absence of mail at 5pm on a Monday in a manner consistent with the carrier's past punctuality. The explanatory power of the conclusion comes, in part, from the fact that if it were a holiday, then there would very probably be no mail. Contrast this with the marbles argument. Though that is a strong argument, the truth of the conclusion 'The next marble I pick will be black', does nothing to explain why 80 of the 100 marbles are black and 20 white. So the marbles argument is not an instance of abductive reasoning—no explanation is attempted.

Keep in mind, as well, that not every explanation is the conclusion of an argument. Hence, not every explanation is part of an abduction. If the statements making up the explanation are already accepted or well known, then no argument is being made, merely an explanation, based in accepted claims.

Returning to the mail carrier, the conclusion presented there is not, however, the only hypothesis which would explain the lack of mail. Flooded roads might well explain the absence of mail. Which is the better, or best, explanation? That depends on how much evidence we have and how much more we can gather. Determining what to count as a good explanation is a complex and interesting philosophical problem, well beyond the scope of this text. We can, however, cite a few useful rules of thumb.

Consider a further example:

> I hear scratching in the walls
> I hear the scurrying-clicking sound of little paws at night
> My cereal and rice boxes have holes chewed in them
> ---
> I have mice

My conclusion, together with other relevant information about the behavior of mice, would well explain the data expressed in the premises. Assuming the presence of mice, the observations are to be expected. Hence, we have a pretty good explanation. Of course, I may have an eccentric neighbor who enjoys practical joking. But until I gather some evidence that she is at work (cheese disappears from the unsprung mousetraps; nary a mouse is to be seen; the sounds of mice do not occur on nights when I lock up the house; there are cheesy fingerprints on the window sill) the mice hypothesis seems to be the best.

Clearly, being subject to inductive criteria, abduction shares all the traits of those criteria. Here are some additional points regarding abduction:

- The more known data an explanation can account for in a consistent and coherent manner, the better the explanation.
- The better an explanation coheres with established theory, the better it is (e.g., hypothesizing that the missing cheese has vanished into thin air, does not (among other problems) cohere well with established physical theory).
- Moreover, if an explanation successfully predicts further data not originally observed (I find footprints outside my garage window; and cheesy fingerprints on the window), then that explanation is better than one which cannot do so.
- There are no precise guidelines for abductive reasoning; especially if we focus on the notion of the 'best' explanation. Much depends on plausibility relative to our background knowledge and the quality and quantity of our evidence. Almost always there is more evidence to be had (in principle), and new evidence (as with inductive, but not valid deductive arguments) can radically alter our assessment of the quality of the inference.

Despite the differences between deductive and inductive criteria, the flowchart for assessing strength and cogency, presented in Figure 2, is parallel to that for validity and soundness (see Fig. 1 on page 10).
It is important to see that the common structure here is the result of the two kinds of question one should ask about any argument. First, the question of the appropriate connection between premises and conclusion (validity or strength): do the premises genuinely support the conclusion? Second, the truth values of the premises (soundness or cogency): are the premises actually true? These are the two points at which any argument may be criticized.

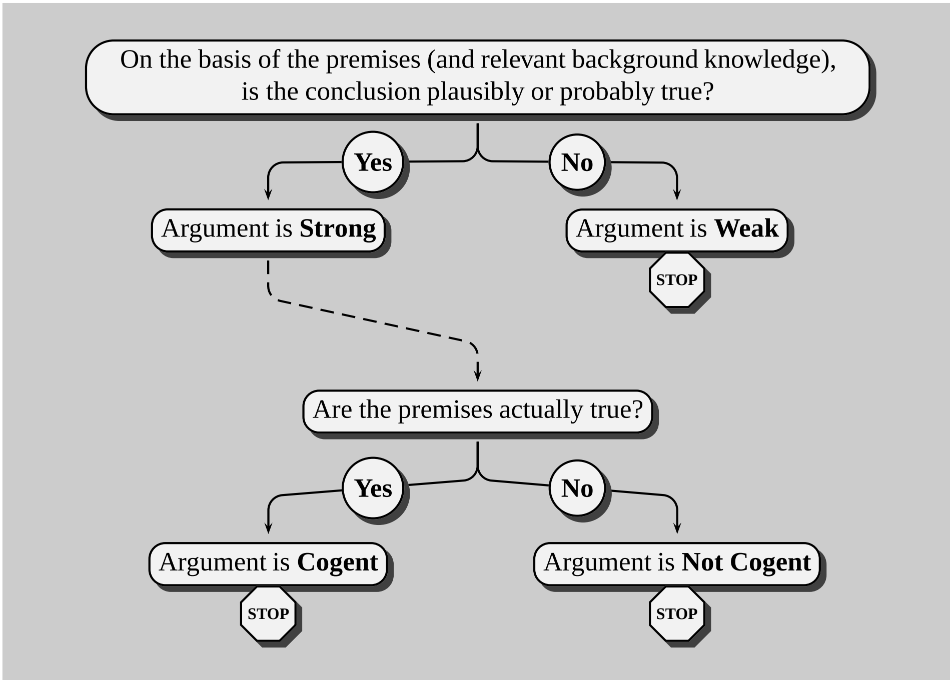

Figure 2: Evaluating Inductive Arguments

6 Fallacies

Generally speaking, a *fallacy* is any mistake in reasoning. But some fallacies are so common that they have earned names.

Fallacy:
A fallacy is any mistake in reasoning, but some are particularly seductive (both to the speaker/writer and the listener/reader) and so common that they have earned names.

Below is a short and incomplete list:

Denying the Antecedent:
See page 11.

Affirming the Consequent:
See page 11.

Arguing in a Circle or Begging the Question:
This involves smuggling one's conclusion (more or less subtly) into one's

premises. For example: "I believe in God because the Bible tells me He exists. I know I can trust the Bible, because it is the received Word of God." But, of course, the speaker is just assuming that God exists (else how could anyone receive His Word?), and using that assumption to argue the same point. Despite the speaker's efforts, he has gone in a circle, and begs the original question: How do you know god exists?

Argument from Authority:

This fallacy involves accepting a belief because someone with some sort of authority has endorsed it. Examples include buying or praising a product simply because a celebrity has endorsed it; or believing in a policy or law simply because your party or some institution with which you are connected endorses it. In some cases, however, we must rely on authority, because we cannot all be experts in every field. We may need to rely on an expert's assessment of scientific or legal claims. However, we should always be ready and willing to challenge an authority intelligently (as opposed to immaturely), and we should be suspicious of "experts" who are unwilling or unable to engage in explanation or justification of their views.

Genetic Fallacy:

This is basically the opposite of the above, occurring when we reject a view simply because of its origins. An example would be disbelieving astronomy because of its origins in astrology; or rejecting a policy out of hand because it is endorsed by the opposing party. While shady origins may be relevant to the current merit of a belief or theory, they are not necessarily relevant, and simply appealing to origins without establishing relevance is not sufficient reason to reject or accept a view.

Ad Hominem (against the man):

This is a specific form of the genetic fallacy in which we attack the character of the view's proponent, as opposed to the reasons offered for the view itself. For example: "Don't believe what Frank says about the CIA, he's just a radical hippie crank." However, there may be cases in which the past behavior of a person is relevant to how reliable or believable they are. If, for instance, Frank can offer no evidence or argument for his position, and has a history of making claims which are later shown to be false, then this is relevant to the question of whether to believe his latest claim. If Frank does offer reasons or evidence, we should critique those, and not his character. If we think character is relevant, then we must establish that relevance if it is to be a basis of criticism.

Straw Man:

We commit this fallacy when we interpret or misrepresent an opponent's position in a way which makes it appear obviously absurd or much weaker than it actually is, then we easily knock down the "straw man" we created. While everyone sometimes says or does absurd things, we should adhere to a principle of charity in interpreting earnest attempts at justifying a view—that is, until faced with clear counterevidence, we should assume that the person is intelligent and is making an honest attempt to explain or justify her claim. Thus, we should try to interpret the person's view in such a way as to give it its strongest and most compelling presentation. If we can criticize that, then so much the better.

Argument from Ignorance:

This fallacy occurs when one argues that because a certain claim is not, or cannot, be proven, then its opposite must be true. "You cannot prove God does not exist, therefore He does." Or, "You cannot prove God does exist, therefore He doesn't." First, outside of mathematics and pure logic, it is unclear whether *any* claim can be "proven"—the best we can do is adduce evidence, reason on its basis, and (in empirical matters) try to test our conclusions. Sometimes this process takes a long time, even decades or centuries. Even given a deductively valid argument in support of a claim, there remains the question of soundness (truth of the premises), and this will require further evidence... Second, this sort of argument ignores the possibility of suspending judgment altogether. Note, though, that there are some cases where repeated failure of active attempts to support or test a theory or claim do weigh against the claim, especially if a better explanation is available.

Equivocation:

This fallacy occurs when the same word or phrase occurs more than once in the premises, but with different meanings. If we do not recognize that two meanings are involved, the argument may appear valid or strong, but when we recognize the equivocation, we see that the conclusion is not supported. For example: "Laws must be made and enforced by someone—a governor or legislative body. Scientists have discovered and stated many laws. Thus, science itself shows there must be a cosmic governor or legislative body—a god or pantheon of gods." In this argument, the word 'law' is being used equivocally. In the first premise, 'law' means a rule of conduct endorsed by a community. In the second premise, 'law' means an invariant regularity in natural phenomena. While the first meaning may imply the existence of a governor or legislative body, the second meaning does not.

False Dilemma:

This occurs when we neglect alternatives or artificially narrow the range of possibilities so as to apparently force the choice of a small range of alternatives. In its simplest form two alternatives are presented, one of which is obviously absurd or undesirable, thus apparently forcing the acceptance of the other. It has the form of a Disjunctive Syllogism (page 9), but it is a false dilemma if alternatives are neglected (i.e., the disjunction is false and the argument unsound). "Either you love your country and obey all its laws, or you are an anti-American radical." But, of course, one can love her country, and still criticize it and try to change its laws—even in radical ways—what could be more American than the willingness to listen to voices of change and potentially agree with them?

This list is not exhaustive, but it gives a good idea of some common mistakes.

One last word about critical thinking. People are often suspicious of logic, seeing it as a sort of "double-speak". But this is usually because they do not know enough logic. Moreover, all of us are to varying degrees resistant to change in our world views—and there are some good reasons for this—but in part this is because we are lazy and don't want to go through the discomfort of having our beliefs shaken up and having to go to the trouble of sorting things out again and possibly (gasp!) changing our minds. In addition to this laziness (which is not a good thing), many people often center their identity—their sense of self—in a certain fixed set of beliefs, or in a certain authoritative institution (such as a church, a political party, an ethnic history, etc.). When you view yourself in this way, any attempt by others (or yourself) to examine or criticize those beliefs or that institution becomes extremely threatening—a form of personal attack—for it threatens the core of your self-identity. A more intelligent, and healthier, alternative is to learn some logic and reasoning, question your own beliefs a bit, and center your self identity not in some set of beliefs or some institution, but in your ability to consider alternatives, to think carefully and critically, and to change your mind if the situation calls for it. This is the difference between being an unreflective believer and a critical thinker. The critical thinker is more flexible and less subject to manipulation by other people and institutions. Hence, the critical thinker is, in an important sense, more free.[6]

[6] Lest I be accused of purveying a false dilemma, I do recognize that these are not the only two ways one can view oneself. In addition, the unreflective believer and the flexible critical thinker are really two poles of a continuum along which our position varies at different times in our lives, even different times of the day, and with regard to different subject matter. My exhortation may then be read as: Be more of a critical thinker more of the time (but don't forget to enjoy yourself)!

7 Exercises

Note: Answers to these exercises appear in Section 9.

1. What is a statement?
2. What is an argument?
3. Explain the meaning of ‘truth value’.
4. Restate the definition of deductive validity.
5. Restate the definition of soundness.
6. Suppose an argument is valid, what, if anything, does this tell us about the truth values of the premises and conclusion?
7. Suppose an argument is sound, what, if anything, does this tell us about whether it is valid?
8. Suppose an argument is sound, what, if anything, does this tell us about the truth values of the premises and conclusion?
9. Suppose an argument has true premises and a true conclusion, what, if anything, does this tell us about the validity and soundness of the argument?
10. Restate the definition of inductive strength.
11. Explain the differences between deductive validity and inductive strength.
12. Explain what is distinctive about abduction.
13. Compose, in standard form, one example of a valid argument and one example of an invalid argument.
14. Using capital letters as placeholders, give the form of the arguments from your previous answer.
15. Compose, in standard form, two instances of the same inductive argument form. Make one instance strong and the other weak. What does this show about the relationship between form and inductive strength?
16. Using the same premises, give two examples of abductive arguments. Make one example strong (a good explanation) and the other weak (a bad explanation). Explain why the one is a better explanation than the other.
17. Find or compose an example of one of the fallacies discussed. Explain why it qualifies as an example of that fallacy.

8 Glossary

Abduction:

Abduction or *abductive reasoning*, also known as *inference to the best explanation*, is a category of reasoning subject to inductive criteria in which the conclusion is supposed to explain the truth of the premises. (20)

Argument Form and Instance:

An *argument form* (or schema) is the framework of an argument which results when certain portions of the component statements are replaced by blanks, schematic letters, or other symbols. An *argument instance* is what results when the blanks in a form are appropriately filled in. (5)

Argument:

An *argument* is a (finite) set of statements, some of which—the *premises*—are supposed to support, or give reasons for, the remaining statement—the *conclusion*. (4)

Cogency:

An argument is *cogent* if and only if it is inductively strong AND all the premises are true. (16)

Conclusion Indicators:

therefore, hence, thus, so, we may infer, consequently, it follows that (5)

Counterexample:

A *counterexample* to an argument (form) is an argument instance of exactly the same form having all true premises and a false conclusion. Production of a counterexample shows that the argument form and all instances thereof are invalid. (Failure to produce a counterexample shows nothing, however). (8)

Deductive Validity, Invalidity:

An argument (form) is *deductively valid* if and only if it is NOT possible for ALL the premises to be true AND the conclusion false. An argument (form) is *deductively invalid* if and only if it is not valid. (6)

Fallacy:

A fallacy is any mistake in reasoning, but some are particularly seductive (both to the speaker/writer and the listener/reader) and so common that they have earned names. (22, see also Section 6)

Inductive Strength:

An argument is *inductively strong* to the degree to which the premises provide evidence to make the truth of the conclusion plausible or probable. If an argument is not strong, it is *weak*. (16)

Logic:

Logic is the study of (i) criteria for distinguishing successful from unsuccessful argument, (ii) methods for applying those criteria, and (iii) related properties of statements such as implication, equivalence, logical truth, consistency, etc. (3)

Logically Consistent:

A set of statements is *logically consistent* if and only if it is possible for all the statements to be true.(14)

Logically Contingent:

A statement is *logically contingent* if and only if it is neither logically true nor logically false; i.e., it is both possible for the statement to be true, and possible for the statement to be false.(12)

Logically Contradictory:

A pair of statements is *logically contradictory* if and only if it is not possible for the statements to have the same truth values.(13)

Logically Entails, Logically Follows:

A set of statements *logically entails* a target statement if and only if it is NOT possible for every member of the set to be true AND the target statement false. We also say that the target statement *logically follows* from the set. (14)

Logically Equivalent:

A pair of statements is *logically equivalent* if and only if it is not possible for the statements to have different truth values.(13)

Logically False:

A statement is *logically false* if and only if it is not possible for the statement to be true.(12)

Logically Inconsistent:

A set of statements is *logically inconsistent* if and only if it is not possible for all the statements to be true.(14)

Logically True:

A statement is *logically true* if and only if it is not possible for the statement to be false.(12)

Premise Indicators:

as, since, for, because, given that, for the reason that, inasmuch as (5)

Soundness:

An argument is *sound* if and only if it is deductively valid AND all its premises are true. (6)

Statement:

A *statement* is a declarative sentence; a sentence which attempts to state a fact—as opposed to a question, a command, an exclamation. (3)

Truth Value:

The *truth value* of a statement is just its truth or falsehood. At this point we make the assumption that every statement is either true (has the truth value true) or false (has the truth value false) but not both. The truth value of a given statement is fixed whether or not we *know* what that truth value is. (4)

9 Answers to §7

1. A statement is a declarative sentence—not a question, command, or exclamation. We assume it is either true or false, but not both.

2. An argument is a set of statements, one of which—the conclusion—is supposed to be supported by the others—the premises.

3. The truth value of a statement is just its truth or falsehood. For our purposes there are two truth values: True and False. Every statement has either the one value or the other, but not both.

4. An argument is deductively valid if and only if it is not possible for all the premises to be true and the conclusion false.

5. An argument is sound if and only if it is valid and all the premises are true.

6. Knowing that an argument is valid does not tell us anything about the actual truth values of the component statements; but it does tell us that if the premises are all true, then the conclusion is true as well.

7. If an argument is sound, then it must be valid as well.

8. If an argument is sound then all the component statements are true. If it is sound then it is valid and has true premises, but this also means that the conclusion is true.

9. Knowing only that all the component statements are true tells us nothing about the validity or soundness of the argument (except that if it is valid, then it is also sound).

10. An argument is inductively strong to the degree to which the premises provide evidence to make the truth of the conclusion plausible or probable.

11. Deductive validity requires that it is not possible for all the premises to be true and the conclusion false. The validity or invalidity of a deductive argument is a matter of form. In inductive arguments there is no attempt at guaranteed truth preservation, but the assumed truth of the premises is supposed to make the conclusion plausible or probable. The strength of an inductive argument is not a matter of form, much background knowledge is relevant to the assessment of strength.

12. Abductive reasoning is a kind of induction in which the conclusion is supposed to explain the premises. Abduction is also known as inference to the best explanation.

13.–17. Check with a classmate or your professor.