

Part I Critical Thinking

Chapter 1 - Basic Training

1.1 Introduction

In this logic course, we are going to be relying on some “mental muscles” that may need some toning up. But don’t worry, we recognize that it may take some time to get these muscles strengthened; so we are going to take it slow before we get up and start running. Are you ready? Well, let’s begin our basic training with some preliminary conceptual stretching. Certainly all of us recognize that activities like riding a bicycle, balancing a checkbook, or playing chess are acquired skills. Further, we know that in order to acquire these skills, training and study are important, but practice is essential. It is also commonplace to acknowledge that these acquired skills can be exercised at various levels of mastery. Those who have studied the piano and who have practiced diligently certainly are able to play better than those of us who have not. The plain fact is that some people play the piano better than others. And the same is true of riding a bicycle, playing chess, and so forth.

Now let’s stretch this agreement about such skills a little further. Consider the activity of thinking. Obviously, we all think, but it is seldom that we think about thinking. So let’s try to stretch our thinking to include thinking about thinking. As we do, and if this does not cause too much of a cramp, we will see that thinking is indeed an activity that has much in common with other acquired skills. As such, it seems reasonable to think that thinking, like other acquired skills, can be improved with training, study and practice. Thinking in this way may be a bit uncomfortable at first, especially if you are inclined to think that thinking is not an acquired skill. Perhaps you believe that thinking is more like a natural function such as eating or drinking, or perhaps like seeing or hearing, than it is like an acquired skill. This may seem plausible, since it is clear that even though not everyone knows how to ride a horse, every normal person eats and drinks, sees and hears, and does so as a matter of course. Accordingly, it would not occur to us to remark on how well one eats her food, as though some are expert eaters, while others are not. Of course, we may need lessons in manners, and perhaps we can all learn to chew our food more than we usually do, but ordinarily we do not think that some people eat better than others, or that some are masters of it, or that others need improvement, or that they need eating lessons, or need further study or more practice in eating. While we may have a good idea of what a chess master is, it seems unlikely that we would as readily know what an eating master would be.

It is therefore tempting to think of thinking as though it were more like eating than like the acquired skills of playing chess or riding a bicycle. Again, while not everyone plays chess or rides horses, everyone eats, and of course everyone thinks. Indeed, if you think of thinking as merely a natural function, as opposed to an acquired skill, you may conclude that you think just fine, or at least that you think as well as anybody else. You may even be offended by the idea that your thinking may be in need of improvement, or in need of further study and practice.

Suppose, however, that thinking is in important respects more like playing a piano or playing chess than it is like eating. Is there any reason to think so? Well, consider this: while it certainly seems to make sense to praise someone as being a good (or bad) thinker, the same does not hold in the case of eating. What would a good (or bad) eater be? This suggests that thinking has something in common with other acquired skills: we can be good or bad at it.

If we can be good at thinking, and not so obviously good at eating (if thinking is more of an acquired skill, like piano playing, than we may have thought at first) then it is reasonable to think that thinking is a skill that, with study and practice, we can improve and even master. Indeed, this is just what we want to suggest. Of course mastering the skill of good thinking will require--as most other such acquired skills do--the mastering of concepts, rules, techniques, and the like. Hard work, however, can be rewarding. I am convinced that if you work hard in this course, if you study hard and most importantly practice and then practice some more, you will improve your thinking skills and become a better and stronger thinker.

Well, what would it be like to become a master thinker? Good question. However, before I try to answer it, I must first alert you to the fact that "thinking well" has various meanings. Such variations occur because there are many different skills available for human mastery. Some people are very good at thinking of new melodies--we call them composers; some people are very good at thinking of novel ways to express human emotions--we call them poets, or novelists; some claim to be better than others at thinking of God, or beauty, or of life in general--we call them mystics. Generally, however, we do not submit these kinds of thinking to appraisals of correctness or incorrectness, or of clarity or a lack of clarity. Moreover, we are not quite sure how we might go about improving these kinds of thinking skills. We take such skills to be natural gifts.

There is, however, another kind of thinking that is subject to being appraised as correct or incorrect. This is the thinking skill called reasoning. To be a good thinker in this sense is to be a person who reasons correctly--most of the time, at least. The master thinker rarely reasons incorrectly.

Even though I want you to think of reasoning well as an acquired skill, I must concede that it is also, in some respects, like a natural function insofar as it does not depend on having special gifts. My assumption, in fact, is that most people reason and that most are capable of reasoning well. In this respect then, thinking well is not like those acquired skills that require special gifts. Although only a few can become accomplished pianists, most can be good thinkers, or at least significantly improve their thinking skills.

Clearly then, we cannot assume that everyone who thinks also reasons well. And unfortunately, as everyone knows, the consequences of bad (incorrect) reasoning are many and may be deep; such consequences range from lost time in solving practical and theoretical problems, to ruined lives caused by coming to unwarranted conclusions about what is true, important, and valuable.

We will begin this course in logic, then, assuming that it makes sense to say that we can improve our reasoning skills, and assuming further that you want to strengthen yours. The way to this improvement will be through study and practice. What you will study are some basic principles and techniques that will enable you, with practice, to distinguish good (correct) reasoning from bad (incorrect) reasoning. With this distinction firmly in mind, you will be able to submit your own reasoning, as well as the reasoning of others, to constructive criticism. This will enable you to listen more carefully to others and to evaluate more critically what they say. It should also make you a more careful and convincing speaker and writer.

So, let's pause and take a breath. I hope that you are beginning to see more clearly what we are going to be doing in this course in logic.

Got your breath? OK, stretching is over. Now it's time to begin our run. It's on to fundamentals.

1.2 Arguments

Well, where have we gotten so far? Far enough, I think, to get to a definition of the subject matter of this course in “logic”. Before we do this, however, let’s make sure we do not confuse logic with psychology. While psychology does study thinking, it does not propose a system for distinguishing good reasoning from bad reasoning.

Keeping this in mind, we can define "logic," the subject matter of this course, as follows:

Logic:
The basic principles and techniques used to distinguish correct (good) reasoning from incorrect (bad) reasoning

If we are to use the term “logic” correctly, however, we must avoid a common confusion. The term is sometimes mistakenly thought to carry only the connotation of good or correct reasoning. For example, when someone has reasoned well (correctly), we say that he or she is being logical, and when someone has reasoned badly (incorrectly) we say that he or she is being illogical. The discipline of logic studies both correct and incorrect reasoning. Accordingly, when someone reasons illogically, we will say that he or she has made a logical mistake, that is, a mistake in logic.

Now, let’s stretch this definition of “**logic**” a little further. As it turns out, **to reason correctly is to argue correctly**. In fact, an alternative definition of the subject matter of logic is as follows: the basic principles and techniques that are used to distinguish **good (correct) arguments** from **bad (incorrect) arguments**. But before we can say more about this distinction between good and bad arguments, we must define what logicians mean by the term “argument.”

“**Argument,**” like so many words, has various meanings. For many it is an entirely negative term connoting confrontation, conflict and disagreement. Clearly, this is what is meant when a coed complains about her date, saying that all she and her boyfriend did was argue all night. The assumption here is that arguing does not make for a pleasant relationship. In addition, most likely, the term “argument” conjures up for some the images of flying pots and pans and shouting, perhaps even violence. Along these lines, we think of arguments as something like schoolyard fights in which the so-called “argument” consists of a heated disagreement in which one party simply denies what the other party affirms, perhaps coming to blows over their difference.

This image of an argument is precisely what is depicted in Monty Python’s famous gag about a person who buys a ticket that entitles him to an argument with a professional arguer. As the gag goes, two people are supposedly arguing, but in fact one person simply denies everything the other person asserts. Such an exchange might go as follows:

“There is life on other planets.”
“No, there's not!” “Oh yes, there is!”
“Oh no, there's not”!
"Is!" "Isn't!" "Is!" "Isn't!"

Back and forth it goes, getting nowhere. While this may be taken for an argument, in the logician’s sense of the term, it is not. We need a more serious definition. So let’s put this caricature of an argument aside, and get on to the logician’s definition.

An argument consists of a group (more than one) of sentences.

Now we must say more precisely what is meant by the word ‘sentences’. There are many kinds of sentences, for example, declarative, imperative, and interrogative sentences just to name a few. The kind of sentences that are used to form arguments are of a specific type. Sentences of this specific type are distinguished from other sentences in having a truth value. Other names for such sentences that have a truth value are often in used by logicians. All of these are virtually synonymous. The most widely used synonymous terms for sentences with a truth value are “propositions” and “statements”. For our purposes, I will use “sentential assertion” instead of “proposition” or “statement.” I note, however, that all of these terms are interchangeable. All of these terms share a common element: they all name the true/false content of some sentential claim. As I will explain presently, in logic we must be careful when we speak of sentences (propositions, statement) to make sure that we are referring only to ones that have a truth value. This is important, as I will explain presently, since many sentences are neither true nor false.

Philosophers often distinguish between two different kinds of sentential claims. One type makes a claim **about existence and the other type makes a claim about meaning**. To keep these two types of sentences distinct, philosophers refer to the former as **synthetic** and the latter as **analytic**. For example, “John is a bachelor” is synthetic; it is about some matter of fact, about something that exists. “Bachelors are not married” is analytic insofar as it simply identifies what bachelors are; that is, it is a remark about what the term ‘bachelors’ means. Analytic sentences are either or false by definition. I also note that a true analytic sentence cannot be false and a false analytic sentence cannot be true. The truth or falsity of analytic sentences is necessary, while the truth or falsity of synthetic sentences is contingent. As well, some sentences that appear to have a truth value do not. Such sentences are neither analytic nor synthetic. These sentences are neither true nor false but nonsense. For example, the sentence “I cannot doubt that I am in pain”, violates what the term ‘pain’ means; it is not false, but nonsense. And one last point, it makes sense to deny a synthetic sentence, but it is senseless to deny an analytic sentence.

While the analytic/synthetic distinction is interesting and important, in this text, our focus will be restricted to sentences that have a truth value. As such, both synthetic and analytic sentences meet this requirement. This does not mean that they both have truth values in the same sense. Synthetic propositions can be **either** true or false, while true analytic sentences cannot be false, and false analytic statements cannot be true. For our purposes, however, we can use either synthetic or analytic sentences to build arguments. That is, more generally, the only sentences that can be used to construct arguments are sentences that have a truth value. Accordingly, I will define an argument as follows:

An argument consists of a group of sentences that have a truth value.

Even though an argument necessarily consists of a group (more than one) of true or false sentences, not every group of sentences is automatically an argument. This is partly because, not all sentences have a truth value. Some sentences are neither true nor false. For example, “I have a pain in my leg but I do not feel it” is not false but senseless. The same can be said for other sentences. Consider the interrogative sentence “Is the cat on the mat?” This sentence is not senseless but it is neither true nor false. And there are others kinds of sentences that are neither true nor false. Consider the exclamatory sentence, “Help!” Clearly, a cry for help is neither true nor false. Other examples abound: “Give me a hamburger”; “I promise”; “I forgive you”; “Whoopee!” These are all sentences, but none are true or false. A group of sentences that are neither true nor false cannot constitute an argument.

And there is another confusion that we must avoid. Do not think that because we do not know whether a sentence is in fact true or false, it is neither. Even though I do not know whether it is raining right now on the White House, the sentence “It is raining right now on the White House” is either true or false. The reason for this is that even if we do not know whether it is or is not in fact raining there, we do know that it is either raining or not raining and not both at the same time. Here we meet with an analytic remark about meaning, not a synthetic claim about existence. Even sentences such as “The universe is finite,” or “God exists” may count as true or false, despite the fact that we may not know whether what is asserted is true or false. What we do know is that what is asserted in these

sentences cannot be both true and false. As I might put this, we know that any given sentential claim is either true or false, and cannot be both, and that this is so even if we do not know what its truth value is in fact. Again, this is a remark about meaning, not existence.

There are a couple more confusions to avoid concerning sentences. Don't think that just because two sentences are exactly the same (contain the same words in the same order) that they therefore express the same content. They may not. How can this be so? Consider the sentence "I have a headache" said by me (Ron Hall), and the same sentence, "I have a headache" said by Jon Smith. The sentences are exactly the same (they have the same words in the same order), but these sentences are used to say different things (and may have different truth values).

It is also possible that different words can be used to express the same sentential assertion. If you will, different words can be used to express the very same thing. Consider the sentence "I have a headache" said by me, and "Ron has a headache" said by Jon. These two sets of words express only one true or false sentence. Now we have different words expressing the same sentence.

And one last thing: one sentence may incorporate multiple sub-sentences. For example, "There must be a fire nearby, because I smell smoke." This is one sentence, but it contains two different sub-sentences: (1) there is a fire nearby and (2) I smell smoke, both of which are either true or false. However, we will say that this is one sentence even though it contains two sub-sentences. Here one sentence is formed by combining two or more sub-sentences. The combination as a whole is then treated as one sentence. "If I smell smoke, then there must be a fire nearby" is one sentence, and has one truth value that applies to the whole sentence, just as each sub-sentence has a separate truth value.

The upshot here is that we need to pay attention to what speakers or writers are doing with sentences. We have to decide whether or not they are saying something that is true or false, that is, determine exactly what they are saying. If we determine that something true or false is being said, then this sentence can be used to build an argument.

An argument must contain a least two sentence that are either true or false. When we have made the determinations that a particular passage (group of sentences) does contain such sentences, it is possible that the passage contains an argument. To make the further determination that the group of sentences constitutes an argument, we must ascertain whether these sentences stand in a special relation to one another. What then is this special relation?

In an argument, true or false sentences are related to each other by virtue of their functions. There are only two such functions: the function of providing support and the function of being supported. In the image below, it is clear that the earth is providing support and the man is being supported.



True or false sentences provide support when they are offered as evidence or as reasons for the truth of another sentence in the group. With this in mind, we can now define an argument as follows:

Argument: A group of true or false sentences expresses an argument only if one or more of these sentences in that group function to support the truth of exactly one other sentence in that group.

Now, let's take this one step further and name these functions that true or false sentences may play in an argument. When these sentences (one or more) that serve the support function in an argument are called **premises**, and the sentence (only one) that is being supported is called the **conclusion** of the argument.

The following table summarizes these definitions:

Premise: A true or false sentence in an argument that functions to support the truth of the conclusion
Conclusion: The sentence in an argument, the truth of which, is supported by the premises
Argument: A group of sentences in which at least one of these functions as a premise and only one of which functions as the conclusion

An argument can have several premises, but only one conclusion. We count arguments by counting conclusions, that is, every argument has only one conclusion.

In addition, there is nothing intrinsic to a particular true or false sentence that makes it either a premise or a conclusion. The sole determination of whether a particular sentence is a premise or a conclusion is its function in the argument. Indeed, the same sentence can function as a conclusion of one argument and a premise in another. Again, if a true or false sentence is functioning in a support role, it is a premise; if it is being supported, then it is a conclusion.

In order for an argument to be persuasive, it must begin with premises that are assumed to be true or probably true. The reason for this is that the premises that are doing the supporting are not themselves supported. Accordingly, they must be accepted if the conclusion of the conclusion is to be accepted. Premises are assumptions. As we might say, the premises of an argument constitute the basis on which the conclusion is established.

The most useful arguments are the ones that arrive at conclusions that were previously not known. Correct useful arguments, in other words, provide us with conclusions that add to our knowledge. This is so because a correct and useful argument establishes or proves that a given sentence (the conclusion) is true or is probably true, depending on the truth of the premises and on the kind of argument it is. (In a moment, we will say more about how conclusions depend on the truth of the premises, and about differences in kinds of arguments.) As well, some arguments do not add to our knowledge. For example, if the conclusion of an argument is analytic, it does not add to our knowledge; it does not inform us of anything; it does not tell us anything about what does or does not exist.

It should be clear by now, that the way logicians think of arguments is quite different from the way that the Monty Python gag represents them. With our expanded concept of an argument, perhaps we are now in a position to understand why someone might rave about his or her date precisely because the two did nothing but argue all night. How exhilarating to think that it is possible for two people actually to have a meaningful and rational conversation. (Looking deeply into each other's eyes over candlelight can last only so long.) So perhaps it makes good sense to imagine the following comment about a date as follows: "What a great date! We argued all evening; she had her opinions, and I had mine, and we were willing to listen to and evaluate the reasons each had for these opinions. This was wonderful!"

Well, I can't promise you a lifetime of such rational conversations, but I hope you will begin to appreciate them and to seek them out. Alternatively, as I might also say, I hope that in this course in logic you will begin to develop your appreciation of the importance and positive value that arguments can have in our lives.

Before you turn to "The Exercised Book", I want to make sure that you have the important distinction between a premise and a conclusion firmly in hand. To do this I remind you to think of it in terms of the difference between **giving and receiving support**. Again, when a sentence is taken as receiving support, it is the conclusion of the argument. When a sentence is used to give support, it is a premise in the argument.

In the image below the man is receiving support from the woman. As we might say, the man is playing the role that a conclusion plays in an argument. As well, the woman is playing the role that a premise plays in an argument.



When you are doing the exercises in the *Exercise Workbook*, keep in mind the definitions that we have just offered. Be especially aware of the logician's special use of terms like "sentence" and "argument." And as I just illustrated, with regard to the logician's special use of the term "argument," be sure that you keep in mind the difference between the role of giving support and the role of receiving support.

1.3 Deduction and Induction

Arguments are usually divided into two primary kinds: inductive arguments (inductions) which are based upon probability and deductive arguments (deductions) which are based upon necessity. We must keep in mind that both are arguments. What is common to both kinds of reasoning is that conclusions are claimed to **follow from premises**. If the conclusions do follow, then the arguments are good ones, and if not, they are bad ones. The distinction between inductions and deductions, then, will turn on two different senses in which conclusions are said to **follow from premises**.

Broadly speaking, an inductive argument offers **some** support for its conclusion, whereas a deductive argument offers support that guarantees the truth of its conclusion. More precisely, we define an inductive argument as follows:

Inductive argument: An argument in which it is claimed that the premises offer some evidence for the conclusion, but do not guarantee the conclusion's truth

Of course the more likely (the more probable) it is that the conclusion follows from the premises, the better the induction. In appraising an inductive argument, we will accordingly say that it is either **weak or strong**— depending on

how likely it is that the truth of the conclusion is established by its premises. It should be clear from this definition that the truth of the premises in an inductive argument is not taken to guarantee the truth of its conclusion. Again, the premises are offered as providing **some evidence** in support of the claim that the conclusion is true. With this definition in mind, it should be obvious that the following argument is an induction:

Every Full Professor in Harvard's Philosophy Department holds a Ph.D. degree. Every Full Professor in Yale's Philosophy Department holds a Ph.D. degree. In fact, this is true of every Philosophy Department with which I am acquainted. Therefore it is likely that every Full Professor in every university Department of Philosophy holds a Ph.D. degree.

This inductive argument is what is called a **generalization**. However, inductive arguments take many forms. These forms include, among others, arguments from analogy, predictions, and sometimes both. Consider the following examples:

1. **Analogy:** *That nest looks like the one I watched a Robin make in my backyard. So it must also be a Robin's nest*
2. **Prediction:** *It usually rains in the late afternoon in the summer when the clouds gather on the horizon like they are doing now. Therefore, I expect it will rain this evening*
3. **Both:** *When I took ancient philosophy in college, we read Plato's Republic. Therefore, if you take that course when you go to college you will read it too.*

Good and bad inductions are so in degree, not absolutely. Additional evidence, of course, can make an inductive argument stronger or weaker. In fact, in some cases one new piece of evidence can completely undermine the likelihood of its conclusion being true. For example, if I am trying to establish that "All swans are white," seeing another white one adds strength to my argument. However, it takes only one black swan to demonstrate that the inductive generalization, "All swans are white," regardless of how many white swans have been observed, is not warranted.

Perhaps this definition of induction is not one that you are familiar with. This is not surprising since there is a rather long tradition of defining induction as reasoning from particular cases to generalizations. In fact the first example of an inductive argument that we cited above fits this definition perfectly. The problem with defining induction this way, however, is that inductive arguments do not always move from particular cases to generalizations. For example, the following argument is an induction, yet it moves from the general to the particular.

Most university professors of philosophy hold the Ph.D. degree. John Smith is a university professor of philosophy. Therefore, John Smith probably holds the Ph.D. degree.

So even though many inductive arguments do move from particular cases to generalizations, not all do. It is therefore a more precise definition of inductive reasoning to say that such arguments claim that the premises of the argument aim at making it likely (probable) that the conclusion is warranted. A strong inductive argument (a good one) succeeds in making the case (by offering premises) that the conclusion is very likely warranted. A weak inductive argument is one that is less successful and, without ceasing to be an induction, may be so weak that it would normally be considered bad reasoning. Consider this very weak induction: "All philosophers are arrogant. I recently had a conversation with two of them at a bar." This is a very weak induction because the pool of evidence expressed in its premise is too small to warrant its conclusion. In an induction, the quantity and quality of the evidence cited in its

premises determines the likelihood or probability of its conclusion and so the strength of the argument. In summary then, we say that **inductive arguments are either weak or strong** and that this assessment is based on the quantity and quality of the evidence cited in its premises.

Now, on to deductive arguments: unlike inductive arguments that offer some evidence for their conclusions, deductive arguments are ones whose premises are offered as a guarantee of the truth of their conclusions. We can define a deductive argument precisely as follows:

Deductive Argument:
An argument in which it is claimed that the conclusion follows from the premises with necessity.

If the conclusion of such an argument does in fact follow with necessity, we say the argument is a good (correct) one and if the conclusion does not follow from the premises with necessity, we say that the argument is a bad (incorrect) one. When we say that the conclusion follows from the premises with necessity in a correct deductive argument, we mean that this conclusion is not just probably true, but must be true, that it cannot fail to be true, **if the premises were true**. When the conclusion does follow necessarily from the premises, we say that this deductive argument is **valid**.

We will turn in the next section to an elaboration of this definition of validity. First, consider the following deductive arguments. I hope that it is obvious to you that these are good arguments.

All rabbits are mammals
All mammals are warm-blooded
Therefore, all rabbits are warm-blooded

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

Given the premises of these arguments, the conclusions are not just likely to be true, but **must be true**, that is, they cannot be false, **if the premises were true**.

Be cautious here: just because we know that an argument is a deduction, we do not know automatically that it is a good one. That is, what makes an argument a deduction is the claim that the truth of the conclusion follows from the truth of the premises with necessity. However, the claim that the conclusion follows from the premises with necessity does not mean that it actually does. In such cases, where the conclusion of a deductive argument may be false even if its premises were true, the argument does not cease to be a deduction; it is simply a bad or incorrect deduction. Consider the following example of an incorrect deductive argument:

All rabbits are mammals
All rabbits have legs
Therefore, all mammals have legs

As in the case of inductions, our definition of deductions differs from the traditional view according to which deductive arguments move from general premises to particular conclusions. Many deductions do indeed follow this pattern. Indeed

the argument above about Socrates is an example of a deductive argument that follows this pattern. However, while the first rabbit argument above is an example of a good deduction and the second rabbit argument that we just cited is an example of a bad deduction, both are cases of deductions that move from general premises to a general conclusion. As well, a deduction can also move from particular premises to a particular conclusion. Consider the following example of such a deduction (which by the way is valid):

*Either Ron Hall is a professor of philosophy or he is an astronaut.
He is not an astronaut. Therefore, Ron Hall is a professor of philosophy.*

So again, we will stay with our more precise definition of deductive arguments, according to which an argument is a deduction if and only if it claims that its conclusion follows from its premises with necessity. While it is important to make the distinction between inductive and deductive arguments, this course will be focused solely on deductive reasoning. Do not take this to mean that I think that deduction is more important than induction. I do not. I recognize, however, that you will have opportunities elsewhere to study induction, since there are many courses offered in statistics and probability. This course in logic, no doubt, will be one of the few opportunities you will have to study deduction as such.

Perhaps the summary will help you to remember the differences between inductive and deductive arguments.

- **Inductive Argument:** Claims that the conclusion is probably true based on the premises
- **Deductive Argument:** Claims that it would be impossible for the conclusion to be false if the premises were true.

INDUCTIVE ARGUMENTS CAN BE:

Strong: The sun has risen in the East for thousands upon thousands of mornings. So it will rise in the East tomorrow morning

Weak: John has a Ph.D. but is unemployed. So this degree does not help you get a job

DEDUCTIVE ARGUMENTS CAN BE:

Correct (Valid): If John practices law in Florida, he must have passed the Florida Bar. John practices law in Florida, so he must have passed the Florida Bar.

Incorrect (Invalid): If John practices law in Florida, he must have passed the Florida Bar. John has passed the Florida Bar. Therefore, John practices law in Florida. (The conclusion here could be false even if the premises were true.)

1.4 Truth, Validity and Soundness

We must be careful not to think that arguments, inductive or deductive, can be evaluated as being either true or false. These terms are reserved solely for sentences. Arguments are either weak or strong, valid or invalid.

Recall that deductive arguments make a claim. That claim is that the conclusion follows from the premises with necessity. Now of course not every claim is warranted. I may claim that I can beat you at tennis. You would be well advised, however, to invite me to the court to see if I can actually make good on this claim.

Similarly, if I claim that a conclusion follows from a set of premises with necessity, this does not mean that the conclusion actually does guarantee the conclusion. We might need to take this claim to a different sort of court, a court of appeals. Well, what are we to appeal to in trying to decide whether or not our claims are satisfied?

Obviously we need a criterion. That is, we need a very precise test for judging whether or not the claim of any deductive argument is satisfied. If a deductive argument passes or fails this test (and the test is strictly pass/fail with nothing in between) we say that the argument is either valid or invalid. So what is this test?

Before I say what the criterion of validity is, and later what the criteria for soundness are, let's note some likely misuses of the terms. Notice that when we introduced the idea of validity above, I did not qualify it as "deductive validity." The reason for this is that the logical term "validity" applies only to deductive arguments.

The logician reserves "validity/invalidity" for evaluating deductive arguments, just as the terms "weak" and "strong" are reserved for evaluating inductive arguments. The same is true in our use of the term "soundness": it is reserved for the evaluation of deductive arguments. That is, inductive arguments cannot be said to be either valid or invalid, nor can they be sound or unsound.

Of course the terms "valid" and "invalid" have uses beyond logic. We have valid or invalid driver's licenses for example. Sometimes people say things. Like: "that is a valid conclusion..." and we know what they mean. But in logic conclusions are either true or false. We can say that a conclusion has been reached validly, in which case we are not talking about its truth but its relation to its premises.

Well enough of these qualifications, we are now ready to make explicit our criterion for validity, and hence invalidity.

Validity:

An argument is valid if and only if it would be impossible for its conclusion to be false if its premises were true, i.e. the truth of the premises guarantees the truth of the conclusion

This criterion is deceptively simple. What is deceptive about it is that it is phrased in the subjunctive mood (which we seldom use anymore), not in the indicative mood, which we almost always use. But this is just a fancy way of saying that our criterion of validity (and hence invalidity) has nothing whatsoever to do with whether or not the premises or conclusion of a particular argument **are** true in fact (indicative mood). Rather, it has everything to do with whether or not the conclusion of a particular argument **would** be true, if its premises **were** true (subjunctive mood).

What is tricky here is to notice that even though the terms "true" and "false" occur in the statement of the criterion of validity, that criterion has nothing to do with what is true or false in fact. Or another way to put this is to say that the judgment of validity (and hence invalidity) is a judgment about the *relation* between the premises and the conclusion of an argument, and **not a judgment about what is true or false in fact.**

Above we talked about conclusions following from premises. Now we are able to give a more precise meaning to this idea. To say that the conclusion of a deductive argument "follows from" the premises is just to say that that conclusion has a certain relation to the premises. Validity is more about this relation than it is about truth. That relation amounts to this:

In valid arguments the conclusion is so related to its premises that it would be impossible for the conclusion to be false if the premises were in fact true, regardless of whether or not they are true in fact.

If you want to see just how little the notion of validity (and hence invalidity) has to do with truth, and how much it has to do with the relation between its premises and its conclusion, consider the following valid argument.

All roaches are millionaires
Ron Hall is a roach
Therefore, Ron Hall is a millionaire

We don't think you will have much trouble agreeing that the first premise of this argument is in fact false. Indeed we hope that you will also concede that the second one is false. But take it from me, Ron Hall is no millionaire. Yet this is a perfectly valid argument. Why? It meets our criterion. **If** the premises **were true**, (and here they are not) then the conclusion **would have to be true** (even though in this case, it is not true in fact).

But think of the situation in which an argument has premises that are true in fact and even a conclusion that is true in fact. Would it have to be valid? No. Consider the following such argument:

All philosophers love wisdom
Jon Smith is a philosopher
Therefore, Jon Smith teaches logic

Now all of these sentences are true in fact, but the argument is not valid. The reason is that the conclusion, even if true in fact, would not have to be true given the truth of the premises. By the same token, if we know that the premises of a given argument are in fact true, and the relation between these premises and its conclusion is valid, we would thereby know that the conclusion would have to be true, that is, could not be false.

As we pointed out above, in a deduction, to argue correctly is to argue validly. And validity is high praise for a deductive argument, even though it does not depend upon the truth of the sentences in the argument. Rather, validity depends solely on the relation between premises and conclusion. Having said this, we are ready to see that validity is not the highest praise we can give to deductive arguments. The reason for this is that truth is important when it comes making the final judgment as to whether or not we have reasoned well. The highest praise for an argument is to say that it is **sound**.

As in the case of validity, we have criteria for soundness (and hence for unsoundness). Those criteria are as follows:

An argument is sound if and only if two conditions are met:

- 1. it is valid**
- 2. all of its premises are true in fact**

Notice that there are two criteria for soundness (and hence for unsoundness). The first criterion is that of validity. However, even though validity is necessary for soundness, and its absence is sufficient for unsoundness, it is not sufficient by itself. Sound arguments must not only be valid, they must also have all true premises. Likewise, having all true premises is not itself a sufficient condition for soundness, even though a single false premise is a sufficient condition for unsoundness. The two criteria, (1) validity and (2) all true premises, taken together are sufficient to determine that a particular argument is sound. Consider the following examples:

1. **Sound Argument:** *If Ron Hall is a philosopher, then Ron Hall loves wisdom. Ron Hall is a philosopher. Therefore, Ron Hall loves wisdom.* (Valid, all true premises)
2. **Unsound Argument A:** *If Jon Smith is a philosopher, then Jon Smith loves wisdom. Jon Smith loves wisdom. Therefore, Jon Smith is a philosopher.* (Invalid, all true premises and a true conclusion)
3. **Unsound Argument B:** *If Socrates is a philosopher then Socrates loves wisdom. Socrates does not love wisdom. Therefore, Socrates is not a philosopher.* (Valid, one false premise and a false conclusion)

Be sure to notice this: even though the criteria of soundness do not mention the conclusion, the conclusion of a sound argument must be true. If an argument is valid, then if its premises were true, its conclusion would necessarily be true. But in sound arguments all the premises are true; hence the conclusion must be true as well.

The upshot of this is that there are two ways to evaluate an argument as being incorrect. We can say that it is invalid or we can say that it is unsound. Just because a conclusion follows from a set of premises, we do not have to accept it as true. It is always open to us to call the truth of one or more of the premises into question.

This issue of truth, however, is not something that we can investigate in a logic course. We will have to leave this to others. In other words, in this course we will not be in the business of evaluating soundness. Rather, we will restrict our evaluations to questions of validity. But this is no small matter, for in order to reason soundly, we must first reason validly. And of course if we reason invalidly, there is no way for our arguments to be sound. The following table represents all of the eight combinations and permutations of truth values that are possible in arguments that have two premises.

- | | |
|---------|--------|
| 1. TTT | 5. FTT |
| 2. TFF | 6. FFF |
| 3. TTF* | 7. FTF |
| 4. TFT | 8. FFT |

In the light of our discussion of validity and soundness ask yourself this question: “Which of these combinations of truth values could represent a valid argument? Careful reflection ought to tell you that every one of these combinations **could** be valid, except one. Do you see which one? Clearly, the combination in case #3 could not ever be valid since it is the only one that directly violates our definition of validity. Remember, if the premises were true then the conclusion would have to be true. In case #3, the premises are true and the conclusion is false. Such a combination could never be valid.

Don’t jump to the conclusion that all of the other combinations are valid. All that we can say is that arguments with any of the other combinations of truth values **could be valid**. You can take any of the other combinations and construct a valid or an invalid argument. For example, we can take combination case 3 and construct both a valid and an invalid argument as follows:

Valid version of case # 4 (TFT): *Either Ron Hall is a philosopher or he likes to walk. He does not like to walk. Therefore he is a philosopher.* (The first premise is true, the second is false, and the conclusion is true.)

Invalid version of case #1 (TTT): *If I go to the movies, I will see Jane. I do not go to the movies. Therefore I do not see Jane.* (If the first and second premises were true and the conclusion happens to be true, the argument is still invalid. This is because the conclusion would not have to be true even if he premises were true. I could have seen Jane at school, in which case the premises could be true and yet the conclusion could be false.)

We could make similar constructions on all of these combinations of truth values except in combination case # 3. In this case, it is impossible to construct a valid argument with this combination of truth values.

Now consider these same combinations in relation to soundness. Ask yourself this question: Which of these combinations could never be sound? Well obviously no sound argument can have any false premise. We can eliminate cases #2, # 3, # 4, # 6, # 7, and #8, since in all of these at least one premise is false. We are left then with cases #1 and #5. Surely, case #5 can be eliminated since it can never be valid, and every sound argument is valid. Therefore, we are left with case #1. Is such a combination always sound? If you reflect carefully you will see that the answer to this question is “no.” Case #1 is the only one of these eight possibilities that **could** be sound. However, the fact remains that even this combination can be invalid. Hence, we can construct an unsound argument with all true premises if the argument we construct is invalid. This is exactly what I just did above in the case of the argument about going to the movies and seeing Jane. Even if all of the premises and the conclusion of an argument are true, the argument will not be sound precisely because it is not valid. It is not valid precisely because its conclusion could be false even though the premises are true in fact.

Ok, its *Workbook* time.

STUDY GUIDE

Chapter 1 - Basic Training

- **Argument:** A group of sentential assertions is an argument when one or more of these sentences in the group is/are used to give evidence (or if you like, reasons, or grounds) in support of one of the other sentences in the group that is receiving support.
- **Conclusion:** A sentential assertion that is the part of an argument that is being supported by the premises.
- **Deductive Argument:** An argument whose premises are claimed to guarantee the truth of its conclusion. Deductive arguments can be either valid or invalid and either sound or unsound.
- **Inductive Argument:** An argument whose premises are claimed to increase the likelihood that its conclusion is true. Inductive arguments can be weak or strong.
- **Logic:** The basic principles and techniques that are used in distinguishing correct (good) reasoning from incorrect (bad) reasoning.
- **Premise:** A sentential assertion that is part of an argument and is used to offer support for the conclusion.
- **Sentence:** A unit of connected speech or writing that can have various uses, including that of expressing true/false claims.
- **Sound Argument:** A valid deductive argument all of whose premises are true.
- **Valid Argument:** A deductive argument whose conclusion would have to be true if its premises were true.