An

Introduction

to

Symbolic Logic

CONTENTS

Introduction

| | *Logic is concerned with arguments, good and bad. With the docile and the reasonable, arguments are sometimes useful in settling disputes. With the reasonable, this utility attaches only to good arguments. It is the logician's business to serve the reasonable. Therefore, in the realm of arguments, it is the logician who distinguishes good from bad.*

                                           Kalish & Montague
1964 p. 1 | |

1 DEDUCTIVE REASONING

Logic is the study of correct reasoning. It is not a study of how this reasoning originates, or what its effects are in persuading people; it is rather a study of what it is that makes some reasoning "correct" as opposed to "incorrect". If you have ever found yourself or someone else making a mistake in reasoning, then this is an example of someone being taken in by incorrect reasoning, and you have some idea of what we mean by correct reasoning: it is reasoning that contains no mistakes, persuasive or otherwise.

It is typical in logic to divide reasoning into two kinds: **deductive** and **inductive**, or, roughly, "airtight" and "merely probable". Here is an example of probable reasoning. You have just been told that Mary bought a new car, and you say to yourself:

> *In the past, Mary always bought big cars.*
> *Big cars are usually gas-guzzlers.*
> *So she (probably) now has a gas-guzzler.*

Your conclusion, that Mary has a gas-guzzler, is not one that you think of as following logically from the information that you have; it is merely a probable inference.

Inductive Logic, which is the study of probable reasoning, is not very well understood at present. There are certain rather special cases that are well developed, such as the application of the probability calculus to gambling games. But a general study has not met with great success. This is not a book about probable reasoning, but if you are interested in it, this is the place to start. This is because most studies of Inductive Logic take for granted that you are already familiar with Deductive Logic -- the logic of "airtight" reasoning -- which forms the subject matter of this book. So you have to start here anyway.

Here is an example of deductive reasoning. Suppose that you recall reading that either James Polk or Eli Whitney was a president of the United States, but you can't remember which one. Some knowledgeable person tells you that Eli Whitney was never president (he was a famous inventor). Based on this information you conclude that Polk was a president.

The information that you have, and the conclusion that you draw from this information, is:

> *Either Polk or Whitney was a president.*
> *Whitney was not a president.*
> *So Polk was a president.*

Let us compare this reasoning with the other reasoning given above. They both have one thing in common: the information that you start with is not known for certain. In the first example, you have only been told that Mary bought a new car, and this may be a lie or a mistake. Likewise, you may be misremembering her past preferences for car sizes. The same is true in the second reasoning: you were only told that Eli Whitney was not president -- by someone else or by a history book -- and your memory that either Polk or Whitney was a president may also be inaccurate. In both cases the information that you start with is not known for certain, and so in this sense your conclusions are only probable. Reasoning is always reasoning from some claims, called the **premises** of the reasoning, to some further claim, called the **conclusion**. If the

premises are not known for certain, then no matter how good the reasoning is, the conclusion will not be known for certain either. (There are certain special exceptions to this; see the exercises below.) There is, however, a difference in the nature of the inferences in the two cases. In the second piece of reasoning the reasoning itself is airtight in the following sense: If the premises that you start with are true, then you are guaranteed that your conclusion is true too. That is, if you were right in thinking that either Polk or Whitney was a president, and if you were right in thinking that Whitney was not a president, then you must be right in thinking that Polk was a president. In this case, it is logically impossible for your premises to be true and your conclusion, nonetheless, to be false. This is an example of what is called **validity**. If your reasoning is valid, then, although you are not guaranteed that your conclusion is true, you are guaranteed that it is true <u>if</u> your premises are.

This guarantee is absent in the case of inductive reasoning. Suppose that Mary has indeed just bought a new car, and suppose that you are correct in believing that she always bought big cars in the past, and also correct in believing that big cars are usually gas-guzzlers. You could still be wrong in your conclusion that she now has a gas-guzzler. Maybe she decided this time to buy a smaller car. Or maybe she got a big one with some extraordinary new fuel economy equipment. These may be unlikely, but they are not ruled out, even assuming that all of your premises are true. The reasoning is not deductively valid because there is a logical possibility that the conclusion is false even if the premises are all true. In short, in the case of inductive reasoning, the inconclusiveness of the reasoning itself introduces further uncertainty in addition to the original uncertainty of the premises.

We rarely have certain knowledge, and a study of logic will not give it to us. Logic is not a method of achieving certainty in general, though it sometimes yields such knowledge as a by-product; instead, it is a study of the logical relationships among all our sentences, including those that are only probable.

2 TRUTH & VALIDITY

A principle unit of investigation in logic is called an **argument**. An "argument", in its technical sense, consists of two parts: a set of sentences, called the **premises**, and a sentence called the **conclusion**. The term "argument" may suggest a dispute, but in logic something is called an argument whether or not any people ever have or ever will disagree about it. Likewise, the "premises" of such an argument may or may not have been believed or asserted by somebody, and it is sometimes useful to examine arguments whose "premises" would never be believed by any rational person. Likewise, by calling something a "conclusion" we do not suggest that anyone ever has or even should "conclude" this thing on the basis of the premises given. The point of the terminology is this: a major topic in the study of deductive logic is **validity**. This is a relationship between a set of sentences and another sentence; this relationship holds whenever it is logically impossible for there to be a situation in which all the sentences in the first set are true and the other sentence false. It turns out to be very useful to study this relationship in complete generality. That is, it is useful to have a theory which tells us when this relationship holds between any set of sentences and any other sentence. Since a major practical application of such a theory is to pieces of reasoning that people actually use, the tradition has arisen of calling the first set of sentences the "premises", and the other sentence the "conclusion". And since a practical application of logic is to situations in which people disagree, it is perhaps appropriate to call the whole thing an "argument". But these are now technical terms. An argument is simply something that has two parts: a set of sentences called the premises, and another sentence called the conclusion. For logical purposes, any such combination counts as an argument.

In displaying arguments it is customary to write their premises first, and to indicate the conclusion by the word like 'so' or a symbol such as '∴'

> *Either Polk was a president or Whitney was a president.*
> *Whitney was not a president.*
> *∴ Polk was a president.*

The triangle made of three dots is an abbreviation of the word `therefore', and is a way of identifying the conclusion of an argument. In order to save on writing, and also to begin displaying the form of the arguments under discussion, we will start abbreviating simple sentences by capital letters. For the time being we will abbreviate *Polk was a president* by `P', and *Whitney was a president* by `W'. We will abbreviate *Whitney was not a president* by `not W'. So the argument can be shortened to:

> P or W
> not W
> ∴ P

A major point of this book is to explore the notion of deductive validity. Since the deductive kind is the only one considered here, we simply refer to it as "validity". In this section we will go over certain consequences of the following definition of validity:

- An argument is **valid** if, and only if, there is no logically possible situation in which all of its premises are true and its conclusion false.

When we talk about "truth" here we do not have anything deep or mysterious in mind. For example, we say that the sentence '*There is beer in the refrigerator*' is true if there is beer in the refrigerator, and false if there isn't beer in the refrigerator. That's all there is to it.

We have already seen one case of a valid argument which has all of its premises true and its conclusion true as well:

> P or W   True
> not W   True
> ∴ P   True

What other possibilities are there? Well, as we noted above, it is possible to have some of the premises false and the conclusion false too. (This is sometimes referred to as a case of the "garbage in, garbage out" principle.) Suppose we use `R' to abbreviate *Robert E. Lee was a president.* Then this argument does not have all of its premises true, nor is its conclusion true:

> R or W   False
> not W   True
> ∴ R   False

Yet this argument is just as good, as far as its validity is concerned, as the first one. If its premises were true, then that would guarantee that its conclusion would be true too. There is no logically possible situation in which the premises are all true and the conclusion false. This argument, though it starts with a false premise and ends up with a false conclusion, has exactly the same **logical form** as the first one. This sameness of logical form lies at the foundation of the theory in this book; it is discussed in the following section.

Although false "inputs" can lead to false outputs, there is no guarantee that this will happen, for you can reason validly from false information and accidentally end up with a conclusion that is true. Here is an example of that:

> P or not W   True
> W   False
> ∴ P   True

In this example, one of the premises is false, but the conclusion happens to be true anyway. Mistaken assumptions can sometimes lead to a true conclusion by chance.

The one combination that we cannot have is a valid argument which has all true premises and a false conclusion. This is in keeping with the definition given above: a deductively valid argument is one for which it is *logically impossible* for its conclusion to be false if its premises are all true.

We have seen that there are valid arguments of each of these sorts:

PREMISES	all true	not all true	not all true
CONCLUSION	true	false	true

Introduction -- 5

What about invalid arguments? (That is, what about arguments that are not deductively valid?)   What combination of truth-values can the parts of invalid arguments have? The answer is that they can have any combination of truth-values whatsoever. Here are some examples:

P or W	True	
P	True	PREMISES ALL TRUE
∴ W	False	CONCLUSION FALSE
P	True	
not W	True	PREMISES ALL TRUE
∴ not R	True	CONCLUSION TRUE
P or W	True	
W	False	PREMISES NOT ALL TRUE
∴ P	True	CONCLUSION TRUE
W or R	False	
P	True	PREMISES NOT ALL TRUE
∴ R	False	CONCLUSION FALSE

The moral of the story so far is that if you know that an argument is invalid, that fact alone tells you nothing at all about the actual truth-values possessed by its parts. And if you know that it is valid, all that that fact tells you about the actual truth-values of its parts is that it does not have all of its premises true plus its conclusion false.

However, there is more to be said. Suppose that you want to show that an argument is invalid, but the argument does not already have all true premises and a false conclusion. How can you do this? One approach is to appeal directly to the characterization of validity, and describe a possible situation in which the premises are all true and the conclusion false. For example, suppose someone has given this (invalid) argument:

> *Either Roosevelt or Truman (or perhaps both) was a president.*
> *Truman was a president.*
> *.: Roosevelt was a president.*

There is no mistake of fact involved here, but the argument is a bad one, and you would like to establish this. You could do so as follows. You say:

> "Suppose that Truman had been a president, but not Roosevelt. In that situation the premises would have been true, but the conclusion false."

This is enough to show the reasoning bad, that is, to show the argument invalid.

We can do even more than this, as we will see in the next section.

EXERCISES

This book provides a stock of exercises as an aid to learning. They were written in the belief that the "hands on" approach to modern logical theory is the best way to master it. You will also be supplied with answers to many of the exercises. You should attempt every exercise on your own, and then check your efforts against the answers that are given. **If you do not understand one or more of the exercises, ask for help!**

Several of the exercises contain material that supplements the explanations in the body of the text. None of the exercises presuppose material that is not provided in the text or in the exercise itself.

1. Decide whether each of the following arguments is valid or invalid. If the argument is invalid then describe a possible situation in which its premises are all true and its conclusion false.

   a. *Either Polk or Lee was a president.*
     Either Lee or Whitney was a president.
     .: Either Polk or Whitney was a president.

   b. *Lee wasn't a president, and Polk was.*
     Either Polk or Whitney was a president.
     .: Whitney was a president.

   c. *Polk was a president and so was Lee.*
     Whitney was a president.
     .: Polk was a president and so was Whitney.

   d. *Either Polk or Whitney was a president.*
     Lee was not a president.
     .: Lee wasn't a president and Polk was.

2. Which of these are true, and which are false:

   a.   Some valid arguments have false conclusions.
   b.   No invalid argument has all true premises and a false conclusion.
   c.   If an argument is valid, and you produce a new argument from it by adding one or more premises to it, the resulting argument will still be valid.
   d.   If an argument is invalid, and you produce a new argument from it by adding one or more premises to it, the resulting argument must still be invalid.
   e.   If an argument has an impossible premise, it is valid. (An example of an impossible sentence is `Some giraffes aren't giraffes'.)
   f.   If an argument has a necessarily true conclusion, it is valid. (An example of a necessarily true sentence is `Every giraffe is a giraffe'.)
   g.   If an argument has a false premise, it is valid.

3. An argument which is valid and which also has all of its premises true is called **sound**. Based on this definition, which of the following are true, and which false:

   a.   All valid arguments are sound.
   b.   All invalid arguments are unsound.
   c.   All sound arguments have true conclusions.
   d.   If an argument is sound, and you produce a new argument from it by adding one or more premises to it, the resulting argument will still be sound.
   e.   All unsound arguments are invalid.
   f.   If an argument has a necessarily true conclusion, it is sound.

4. Suppose that we have two arguments which are both valid:

        W   Q
       ∴ Q   ∴ S

what do we know about this argument?

        W
    ∴ S

5. Suppose that the following two simple arguments are both sound:

        W   Q
    ∴ Q   ∴ S

what do we know about this argument?

        W
    ∴ S

6. Suppose that the following argument is invalid:

        W
    ∴ S

what do we know about these arguments?

        W   R
    ∴ R   ∴ S

7. (a) Give an example of a "reversing" argument, that is, one which is guaranteed to have a false conclusion if its premises are true, and is guaranteed to have a true conclusion if any of its premises are false. (b) Give an example of an argument that must have a true conclusion no matter what the truth-values of its premises. Is this argument valid?

3 LOGICAL FORM

If you want to show that an argument is invalid, you can describe a possible situation in which the premises are all true and the conclusion false. But this direct appeal to possible situations is sometimes difficult to articulate, and judgments of possibility can differ. Fortunately, there is another technique that is often more useful. You could have challenged the above reasoning by saying:

> "That reasoning is no good. If that reasoning were good,
> we could prove that McGovern was a president! For we know that:

> *Either McGovern or Truman was a president,*

> and we know:

> *Truman was a president.*

> So, by your reasoning we should be able to conclude that

> *McGovern was a president too*!"

This challenge, like the first one, also shows that the argument given above is invalid. But whereas the first type of challenge focuses on how the ORIGINAL argument works in some POSSIBLE situation, this second challenge is based on how some OTHER argument works in the ACTUAL situation. What we do in this second technique is to give an argument that is different than the first, but closely related to it. In the case in question, the new argument is:

> *Either McGovern or Truman was a president.*
> *Truman was a president.*
> *∴ McGovern was a president.*

We know the new argument is invalid because it actually has all true premises and a false conclusion (we chose it on purpose to be this way). Since the new argument is invalid, so is the original one.

But why should the original argument be invalid just because this second argument is invalid? The answer is that, intuitively speaking, they both employ the same reasoning, and it is the reasoning that is being assessed when we make a judgment about validity. But how can we tell that they employ the same reasoning? The answer is that they both have the **same form**. Each argument is one in which one of the premises is an "or" statement, with the other premise being one of the parts of the "or" statement and the conclusion being the other part. This sameness of structure or form indicates a sameness of the reasoning involved.

A key assumption on which all of modern logical theory is based is that goodness of deductive reasoning is a matter of **form**, not a matter of content. Any argument which has just the same form as the argument we were just discussing is invalid, no matter whether its subject matter is religion, politics, mathematics, or baseball. Likewise, any argument which has this form:

> P or W
> not W
> ∴ P

is valid, regardless of its subject matter. With this in mind, we can give a modern account of validity due to form:

- An argument is **formally valid** if and only if every argument with exactly the same form is valid.

It follows from this definition that if an argument is formally valid, so is any argument with exactly that form, if an argument is not formally valid, neither is any argument with exactly that form. A central preoccupation of modern logic, then, is the investigation and classification of logical forms. (That is why this logic is called "*formal* logic".) This will be our business throughout the chapters that follow.

EXERCISES

1. Decide whether each of the following arguments is valid or invalid. If the argument is invalid then give an argument which has the same form, and which actually has all true premises and a false conclusion.

   a. *Either Polk or Lee was a president.*
     Either Whitney or Lee was a president.
     ∴ Either Polk or Whitney was a president.

   b. *Lee wasn't a president, and Polk was.*
     Either Polk or Whitney was a president.
     ∴ Whitney was a president.

   c. *Polk was a president and so was Lee.*
     Whitney was a president.
     ∴ Polk was a president and so was Whitney.

   d. *Either Polk or Whitney was a president.*
     Lee was not a president.
     ∴ Lee wasn't a president and Polk was.

2. Which of these are true, and which are false:

  a. Some invalid arguments have the same forms as valid ones.
  b. You can show an argument valid by producing another argument which has the same form and which has true premises and a true conclusion.

  c. If you are wondering whether an argument is valid or not, and you fail to find another argument which has the same form and all true premises and a false conclusion, that shows the original argument to be valid.

3. Here are some argument forms. For each, say whether every argument with that form is valid. If it is not valid, give an example of an argument with the given form that has true premises and a false conclusion.

   a.   If A then B
                    A
                    ∴ B

   b.   If A then B
                    B
                    ∴ A

   c.   not (A and B)
                    not-B
                    ∴ not A

   d.   A or B
                    B
                    ∴ A

   e.   A and not-A
                    ∴ B

   f.   A
                    ∴ B or not-B

   g.   A or B
                    not-B or C
                    ∴ A or C

4. Recall that an argument which is valid and which also has all of its premises true is called sound.

   a.   If you are wondering whether an argument is sound, and you manage to find another one with the same form and having all true premises and a false conclusion, does that show the original argument to be unsound? Why?

   b.   If you are wondering whether an argument is sound, and you manage to find another one with the same form and having all true premises and a true conclusion, does that show the original argument to be sound? Why?

   c.   If you are wondering whether an argument is sound, and you manage to find another one with the same form and having all false premises and a false conclusion, does that show the original argument to be sound? To be unsound? Why?

5. For each of the examples in 3, say whether or not *every* argument with that form is sound, and also say whether *some* argument with that form is sound.

6. {This question is speculative, and does not necessarily have a straightforward answer} Could there be an argument that is valid but not formally valid? Could there be an argument that is formally valid but not valid?

4 SYMBOLIC NOTATION

Our investigation of logical forms will take an indirect route, but one that has proved to be worthwhile. Instead of attempting a direct classification of the logical forms of sentences of English, we will develop an artificial language that is considerably simpler than English. It will in some ways be like English without the logically irrelevant aspects of English. And it will lack some of the characteristics that make the use of English confusing when used in argumentation. For example, the artificial language will lack some of the structural ambiguity of English. Consider this English sentence:

> *Mary teaches little girls and boys.*

Does this tell us that Mary teaches little girls and little boys, or that she teaches little girls and regular-size boys? If this sentence occurred in an argument, the validity of the argument might turn on how the sentence was read. In the artificial language to be developed, structural ambiguities of this sort will be absent.

The artificial language will be especially designed to make logical form perspicuous. You are already familiar with this from arithmetic. Consider the partly symbolic sentence:

> *For any two numbers x and y, x+y = y+x.*

It is clear what this says. The same thing can also be said without any symbols:

> *Given any two numbers, the result of adding them together in one order is the same as the result of adding them together in the reverse order.*

It is apparent that the use of symbols makes the claim clearly and vividly. Our logical symbolism will be like this.

In fact, it will be possible in the symbolic language to tell the logical forms of its sentences just by examining the shapes and arrangements of its symbols. And it will be possible to evaluate formal validity and invalidity of arguments expressed in it by a variety of techniques that appeal directly to the visible arrangements of the symbols in the sentences used.

This does not mean that we will lose sight of reasoning that is expressed in our native tongue. One of the tasks of learning the artificial language will be to learn how to take sentences of English and re-express them in the artificial language we are learning. One of our goals, after all, is to learn about reasoning that we encounter every day, in its natural habitat.

5 IDEALIZATIONS

The data that we have to deal with are incredibly complex, and this is only an introductory text. So we will idealize from time to time. This is no different from any other art or science. In physics you usually begin by studying the behavior of bodies falling in a perfectly uniform gravitational field, or sliding down frictionless planes. There are no perfectly uniform gravitational fields, and no frictionless planes, yet studying these things gives us clear and simple models that can be applied to real phenomena as approximations. And then in the advanced courses you can learn how friction affects the sliding, and how non-uniform fields affect the movement of things in them.

Here are some of the idealizations that we will make in this book: We will look only at arguments with indicative sentences, not with imperative or interrogative. We will ignore any problems due to vagueness. For example, given a perfect understanding of the situation, you may still be unsure whether to say that *Mary loves John*, because of the vagueness of distinguishing between loving and liking. We will also totally ignore the fact that sentences may change truth-value over time and with differing situations. If I say today:

> *I'm feeling great!*

this may be true, but the very same sentence may be false tomorrow. And it may be true when I say it, yet false when someone else utters it. This "context dependence" of truth has aroused a great deal of interest, and there are many theories about how it works. They all presuppose that their readers have already learned the material in this book. We will pretend in our investigations that sentences come with unique truth-values that do not change with context. The effects of context constitute an advanced study.

We will also assume that each sentence is either true or false. Again, the question of whether, and which, sentences lack truth-value is interesting, but is not to be pursued at the beginning.

Many other idealizations will become apparent as we proceed.

6 THE PLAN OF THE TEXT

In this text we will develop a symbolic notation in a step-by-step process. In chapter 1 we consider simple sentences that are combined together with the "connectives" '*or*' and '*if . . .then . . .*'. Even with this very austere notation we can formulate and study a number of formal validities. In chapter 2 we expand this notation by introducing further connectives: '*and*', '*or*', and '*if and only if*'. In chapter 3 we vastly expand our symbolic language with the introduction of variables and quantifiers. Each expansion of the notation builds on what has gone before, so we are continually increasing our ability to validate formally valid arguments and invalidate arguments that are not formally valid.