CONTENTS

1

Propositions and Arguments

1.1 What Is Logic?

Somebody who wants to do a good job of measuring up a room for purposes of cutting and laying carpet needs to know some basic mathematics—but mathematics is not the science of room measuring or carpet cutting. In mathematics one talks about angles, lengths, areas, and so on, and one discusses the laws governing them: if this length is smaller than that one, then that angle must be bigger than this one, and so on. Walls and carpets are things that have lengths and areas, so knowing the general laws governing the latter is helpful when it comes to specific tasks such as cutting a roll of carpet in such a way as to minimize the number of cuts and amount of waste. Yet although knowing basic mathematics is essential to being able to measure carpets well, mathematics is not rightly seen as the science of carpet measuring. Rather, mathematics is an abstract science which gets applied to problems about carpet. While mathematics does indeed tell us deeply useful things about how to cut carpets, telling us these things is not essential to it: from the point of view of mathematics, it is enough that there be angles, lengths, and areas considered in the abstract; it does not matter if there are no carpets or floors.

Logic is often described as the study of *reasoning*.[1] Knowing basic logic is indeed essential to being able to reason well—yet it would be misleading to say that human reasoning is the primary subject matter of logic. Rather, logic stands to reasoning as mathematics stands to carpet cutting. Suppose you are looking for your keys, and you know they are either in your pocket, on the table, in the drawer, or in the car. You have checked the first three and the keys aren't there, so you reason that they must be in the car. This is a good way to reason. Why? Because reasoning this way cannot lead from true premises or starting points to a false conclusion or end point. As Charles Peirce put it in the nineteenth century, when modern logic was being developed:

> The object of reasoning is to find out, from the consideration of what we already know, something else which we do not know. Consequently, reasoning is good if it

> be such as to give a true conclusion from true premises, and not otherwise. [Peirce, 1877, para. 365]

This is where logic comes in. Logic concerns itself with *propositions*—things that are true or false—and their components, and it seeks to discover laws governing the relationships between the truth or falsity of different propositions. One such law is that if a proposition offers a fixed number of alternatives (e.g., the keys are either (i) in your pocket, (ii) on the table, (iii) in the drawer, or (iv) in the car), and all but one of them are false, then the overall proposition cannot be true unless the remaining alternative is true. Such *general* laws about truth can usefully be applied in reasoning: it is because the general law holds that the particular piece of reasoning we imagined above is a good one. The law tells us that if the keys really are in one of the four spots, and are not in any of the first three, then they must be in the fourth; hence the reasoning cannot lead from a true starting point to a false conclusion.

Nevertheless, this does not mean that logic is itself the science of reasoning. Rather, logic is the science of *truth.* (Note that by "science" we mean simply *systematic study.*)[2] As Gottlob Frege, one of the pioneers of modern logic, put it:

> Just as "beautiful" points the ways for aesthetics and "good" for ethics, so do words like "true" for logic. All sciences have truth as their goal; but logic is also concerned with it in a quite different way: logic has much the same relation to truth as physics has to weight or heat. To discover truths is the task of all sciences; it falls to logic to discern the laws of truth. [Frege, 1918–19, 351]

One of the goals of a baker is to produce hot things (freshly baked loaves). It is not the goal of a baker to develop a full understanding of the laws of heat: that is the goal of the physicist. Similarly, the physicist wants to produce true things (true theories about the world)—but it is not the goal of physics to develop a full understanding of the laws of truth. That is the goal of the logician. The task in logic is to develop a framework in which we can give a detailed—yet fully general—representation of propositions (i.e., those things which are true or false) and their components, and identify the general laws governing the ways in which truth distributes itself across them.

Logic, then, is primarily concerned with truth, not with reasoning. Yet logic is very usefully applied to reasoning—for we want to avoid reasoning in ways that could lead us from true starting points to false conclusions. Furthermore, just as mathematics can be applied to many other things besides carpet cutting, logic can also be applied to many other things apart from human reasoning. For example, logic plays a fundamental role in computer science and computing technology, it has important applications to the study of natural and artificial languages, and it plays a central role in the theoretical foundations of mathematics itself.

1.2 Propositions

We said that logic is concerned with the laws of truth. Our primary objects of study in logic will therefore be those things which can be true or false—and so it will be convenient for us to have a word for such entities. We shall use the term "proposition" for this purpose. That is, *propositions* are those things which can be true or false. Now what sort of things are propositions, and what is involved in a proposition's being true or false? The fundamental idea is this: a proposition is a claim about how things are—it represents the world as being some way; it is true if the world is that way, and otherwise it is false. This idea goes back at least as far as Plato and Aristotle:

> SOCRATES: But how about truth, then? You would acknowledge that there is in words a true and a false?
> HERMOGENES: Certainly.
> S: And there are true and false propositions?
> H: To be sure.
> S: And a true proposition says that which is, and a false proposition says that which is not?
> H: Yes, what other answer is possible? [Plato, c. 360 BC]

> We define what the true and the false are. To say of what is that it is not, or of what is not that it is, is false, while to say of what is that it is, and of what is not that it is not, is true. [Aristotle, c. 350 BC-a, Book IV (Γ) §7]

In contrast, nonpropositions do not represent the world as being thus or so: they are not claims about how things are. Hence, nonpropositions cannot be said to be true or false. It cannot be said that the world is (or is not) the way a nonproposition represents it to be, because nonpropositions are not claims that the world is some way.[3]

Here are some examples of propositions:

1. Snow is white.

2. The piano is a multistringed instrument.

3. Snow is green.

4. Oranges are orange.

5. The highest speed reached by any polar bear on 11 January 2004 was 31.35 kilometers per hour.

6. I am hungry.

Note from these examples that a proposition need not be true (3), that a proposition might be so obviously true that we should never bother saying it was true (4), and that we might have no way of knowing whether a proposition

is true or false (5). What these examples *do* all have in common is that they make claims about how things are: they represent the world as being some way. Therefore, it makes sense to speak of each of them as being true (i.e., the world is the way the proposition represents it to be) or false (things aren't that way)—even if we have no way of knowing which way things actually are.

Examples of nonpropositions include:

7. Ouch!

8. Stop it!

9. Hello.

10. Where are we?

11. Open the door!

12. Is the door open?

It might be appropriate or inappropriate in various ways to say "hello" (or "open the door!" etc.) in various situations—but doing so generally could not be said to be true or false. That is because when I say "hello," I do not make a claim about how the world is: I do not represent things as being thus or so.[4] Nonpropositions can be further subdivided into *questions* (10, 12), *commands* (8, 11), *exclamations* (7, 9), and so on. For our purposes these further classifications will not be important, as all nonpropositions lie outside our area of interest: they cannot be said to be true or false and hence lie outside the domain of the laws of truth.

1.2.1 Exercises

Classify the following as propositions or nonpropositions.

1. Los Angeles is a long way from New York.

2. Let's go to Los Angeles!

3. Los Angeles, whoopee!

4. Would that Los Angeles were not so far away.

5. I really wish Los Angeles were nearer to New York.

6. I think we should go to Los Angeles.

7. I hate Los Angeles.

8. Los Angeles is great!

9. If only Los Angeles were closer.

10. Go to Los Angeles!

1.2.2 Sentences, Contexts, and Propositions[5]

In the previous section we stated "here are some examples of propositions," followed by a list of sentences. We need to be more precise about this. The

idea is not that each sentence (e.g., "I am hungry") is a proposition. Rather, the idea is that *what the sentence says* when uttered in a certain context—the claim it makes about the world—is a proposition.[6] To make this distinction clear, we first need to clarify the notion of a *sentence*—and to do that, we need to clarify the notion of a *word:* in particular, we need to explain the distinction between word *types* and word *tokens.*[7]

Consider a word, say, "leisure." Write it twice on a slip of paper, like so:

leisure	leisure

How many words are there on the paper? There are two word tokens on the paper, but only one word type is represented thereon, for both tokens are of the same type. A word token is a physical thing: a string of ink marks (a flat sculpture of pigments on the surface of the paper), a blast of sound waves, a string of pencil marks, chalk marks on a blackboard, an arrangement of paint molecules, a pattern of illuminated pixels on a computer screen—and so on, for all the other ways in which words can be physically reproduced, whether in visual, aural, or some other form. A word token has a location in space and time: a size and a duration (i.e., a lifespan: the period from when it comes into existence to when it goes out of existence). It is a physical object embedded in a wider physical context. A word type, in contrast, is an abstract object: it has no location in space or time—no size and no duration. Its *instances*—word tokens—each have a particular length, but the word type itself does not. (Tokens of the word type "leisure" on microfilm are very small; tokens on billboards are very large. The word type itself has no size.) Suppose that a teacher asks her pupils to take their pencils and write a word in their notebooks. She then looks at their notebooks and makes the following remarks:

1. Alice's word is smudged.

2. Bob and Carol wrote the same word.

3. Dave's word is in ink, not pencil.

4. Edwina's word is archaic.

In remark (1) "word" refers to the word token in Alice's book. The teacher is saying that this token is smudged, not that the word type of which it is a token is smudged (which would make no sense). In remark (2) "word" refers to the word type of which Bob and Carol both produced tokens in their books. The teacher is not saying that Bob and Carol collaborated in producing a single word token between them (say by writing one letter each until it was finished); she is saying that the two tokens that they produced are tokens of the one word type. In remark (3) "word" refers to the word token in Dave's book. The teacher is saying that this token is made of ink, not that the word type of which

it is a token is made of ink (which, again, would make no sense). In remark (4) "word" refers to the word type of which Edwina produced a token in her book. The teacher is not saying that Edwina cut her word token from an old manuscript and pasted it into her book; she is saying that the word type of which Edwina produced a token is no longer in common use.

Turning from words to sentences, we can make an analogous distinction between sentence types and sentence tokens. Sentence types are abstract objects: they have no size, no location in space or time. Their instances—sentence tokens—do have sizes and locations. They are physical objects, embedded in physical contexts: arrangements of ink, bursts of sound waves, and so on. A sentence type is made up of word types in a certain sequence;[8] its tokens are made up of tokens of those word types, arranged in corresponding order. If I say that the first sentence of Captain Cook's log entry for 5 June 1768 covered one and a half pages of his logbook, I am talking about a sentence token. If I say that the third sentence of his log entry for 8 June is the very same sentence as the second sentence of his log entry for 9 June, I am talking about a sentence type (I am not saying of a particular sentence token that it figures in two separate log entries, because, e.g., he was writing on paper that was twisted and spliced in such a way that when we read the log, we read a certain sentence token once, and then later come to that very same token again).[9]

§

Now let us return to the distinction between sentences and propositions. Consider a sentence type (e.g., "I am hungry"). A speaker can make a claim about the world by uttering this sentence in a particular context. Doing so will involve producing a token of the sentence.[10] We do not wish to identify the proposition expressed—the claim about the world—with either the sentence type or this sentence token, for the reasons discussed below.

To begin, consider the following dialogue:

Alan: Lunch is ready. Who's hungry?
Bob: I'm hungry.
Carol: I'm hungry.
Dave: I'm not.

Bob and Carol produce different tokens (one each) of the same sentence type. They thereby make different claims about the world. Bob says that he is hungry; Carol says that she is hungry. What it takes for Bob's claim to be true is that Bob is hungry; what it takes for Carol's claim to be true is that Carol is hungry. So while Bob and Carol both utter the same sentence type ("I'm hungry") and both thereby express propositions (claims about the world), they do not express the same proposition. We can be sure that they express different propositions, because what Bob says could be true while what Carol says

is false—if the world were such that Bob was hungry but Carol was not—or vice versa—if the world were such that Carol was hungry but Bob was not. It is a sure sign that we have two distinct propositions—as opposed to the same proposition expressed twice over—if there is a way things could be that would render one of them true and the other false.[11] So one sentence type can be used to express different propositions, depending on the context of utterance. Therefore, we cannot, in general, identify propositions with sentence types.[12]

Can we identify propositions with sentence tokens? That is, if a speaker makes a claim about the world by producing a sentence token in a particular context, can we identify the claim made—the proposition expressed—with that sentence token? We cannot. Suppose that Carol says "Bob is hungry," and Dave also says "Bob is hungry." They produce two different sentence tokens (one each); but (it seems obvious) they make the same claim about the world. Two different sentence tokens, one proposition: so we cannot identify the proposition with both sentence tokens.[13] We could identify it with just one of the tokens—say, Carol's—but this would be arbitrary, and it would also have the strange consequence that the claim Dave makes about the world is a burst of sound waves emanating from Carol. Thus, we cannot happily identify propositions with sentence tokens.

Let us recapitulate. A proposition is a claim about how things are: it represents the world as being some way. It is true if things are the way it represents them to be (saying it how it is) and otherwise it is false (saying it how it isn't). The main way in which we make claims about the world—that is, express propositions—is by uttering sentences in contexts. Nevertheless, we do not wish to identify propositions with sentences (types or tokens), because of the following observations:

- One sentence type can be used (in different contexts) to make distinct claims (the example of "I'm hungry," as said by Bob and Carol).

- The same claim can be made using distinct sentence types (the example of John and Johann's sentences in n. 12).

- The same claim can be made using distinct sentence tokens (the example of Carol's and Dave's tokens of "Bob is hungry").

It should be said that we have not discussed these issues in full detail.[14] We have, however, said enough to serve our present purposes. For we do not wish to deny vehemently and absolutely that propositions might (in the final analysis) turn out to be sentences of some sort. Rather, we simply wish to proceed without assuming that propositions—our main objects of study— can be reduced to something more familiar, such as sentences. In light of the problems involved in identifying propositions with sentences, our decision to refrain from making any such identification is well motivated.

So far so good, then. But now, if propositions are not sentences, then what are they? Propositions might start to seem rather mysterious entities. I can picture tokens of the sentence "I am hungry," and perhaps, in some sense, I can even picture the sentence type (even though it is an abstract object). But how do I picture the proposition that this sentence expresses (when a certain speaker utters it in a particular context)? It would be a mistake in methodology to try to answer this question in detail at this point. One of the tasks of logic—the science of truth—is to give us an understanding of propositions (the things that are true or false). What we need in advance of our study of logic—that is, what we need at the present point in this book—is a rough idea of what it is of which we are seeking a precise account. (Such a rough idea is needed to guide our search.) But we now have a rough idea of what propositions are: they are claims about the world; they are true if the world is as claimed and otherwise false; they are expressed by sentences uttered in context but are not identical with sentence types or with tokens thereof. The detailed positive account of propositions will come later (§11.4).

§

There is one more issue to be discussed before we close this section. We have seen that, to determine a proposition, we typically need not just a sentence type but also a context in which that sentence is uttered. For example, for the sentence type "I am hungry" to determine a proposition, it needs to be uttered by someone in a particular context. When Bob utters it, he then expresses the proposition that he (Bob) is hungry (that is how the world has to be for what he says to be true); when Carol utters it, she then expresses the proposition that she (Carol) is hungry (that is how the world has to be for what she says to be true); and so on. This general picture is widely accepted. However, exactly *how* it comes about that a particular proposition is expressed by uttering a certain sentence in a specific context is a topic of great controversy. Some of the factors that potentially play a role in this process are:

1. The meaning of the sentence type. (This is usually thought of as determined by the meanings of the sentence's component word types together with the syntax—the grammatical structure—of the sentence. The meaning of a word type is what a speaker has to know to understand that word; it is what a dictionary entry for that word aims to capture.)

2. Facts about the context of utterance. (Relevant facts include the time and place of the context, and the identity of the speaker.)

3. Facts about the speaker (e.g., what she intended to convey by uttering the sentence she uttered in the context).

Together, these facts—and perhaps more besides—determine what is said by the speaker in uttering a certain sentence in a certain context; that is, what claim she is making about the world—which proposition is expressed.[15] That much is widely agreed; the controversy enters when it comes to the question of exactly how the labor of determining a particular proposition is divided between the contributing factors mentioned above: what role each plays. We do not need to enter these controversies here, however: for in logic we are concerned with propositions themselves, not with how exactly they come to be expressed by uttering sentences in contexts.[16] This is not to say that in this book we shall be able to get by without sentences. On the contrary, our chief way of getting in touch with propositions is via the sentences that express them. The point to keep in mind is that our primary interest is in the propositions expressed: sentences are simply a route to these propositions.[17]

1.3 Arguments

We said that the laws of truth underwrite principles of good reasoning. Reasoning comes packaged in all sorts of different forms in ordinary speech, writing, and thought. To facilitate discussion of reasoning, it will be useful to introduce a *standard form* in which any ordinary piece of reasoning can be represented. For this purpose we introduce the notion of an *argument*. As was the case with the term "proposition," our usage of the term "argument" is a technical one that is abstracted from one of the ordinary meanings of the term. In our usage, an argument is a *sequence of propositions*. We call the last proposition in the argument the *conclusion:* intuitively, we think of it as the claim that we are trying to establish as true through our process of reasoning. The other propositions are *premises:* intuitively, we think of them as the basis on which we try to establish the conclusion. There may be any finite number of premises (even zero). We may present arguments in the following format:

> Premise 1
> Premise 2
> ————————
> Conclusion

Here we use a horizontal line to separate the conclusion from the premises. The conclusion can also be marked by the term "Therefore" (often abbreviated as ∴):

> Premise 1
> Premise 2
> Premise 3
> ∴ Conclusion

We may also present an argument in a linear fashion, with the premises separated by commas and the conclusion separated by a slash and the "therefore" symbol:

Premise 1, Premise 2, Premise 3, Premise 4 /∴ Conclusion

For example, consider the following piece of ordinary reasoning. I do not have a watch, and I am wondering what time it is. I notice that Sesame Street is just starting on television, and I know from my acquaintance with the timetable that this show starts at 8.30. I conclude that it is now 8.30. We can represent this piece of reasoning as the following argument:

If Sesame Street is starting, it is 8.30.
Sesame Street is starting.

∴ It is 8.30.

When looking at a piece of reasoning phrased in ordinary language with a view to representing it as an argument, we identify the conclusion as the proposition that the speaker is trying to establish—to give reasons for—and the premises as the reasons given in support of that conclusion. Phrases that commonly indicate conclusions in ordinary reasoning include "therefore," "hence," "thus," "so," and "it follows that;" phrases that commonly indicate premises include "because," "since," and "given that." However, these words are not always present, and even when they are they do not always indicate conclusions and premises, respectively. Hence there is no recipe we can follow mechanically when representing ordinary reasoning in the form of an argument: we must always think carefully about what is being said in the ordinary reasoning—about what it is that the reasoner is trying to establish (this will be the conclusion) and about what reasons are being given in support of this conclusion (these will be the premises). One point to note carefully is that when we represent a piece of reasoning as an argument in our technical sense—that is, a sequence of propositions—we always put the conclusion last. In ordinary English, however, the conclusion of a piece of reasoning is not always what is stated last.

Let's consider another example. When working out what to serve a guest for breakfast, someone might reason as follows: Mary must like marmalade, because she is English, and all English people like marmalade. Here the conclusion—the proposition that the reasoning is supposed to establish—is the thing said first: that Mary likes marmalade. The premises are the reasons given in support of this conclusion—that Mary is English, and that all English people like marmalade. So we represent this piece of reasoning as the following argument:

Mary is English.
All English people like marmalade.
Therefore, Mary likes marmalade.

Note that we count any sequence of one or more propositions as an argument. Thus we count as arguments things that do not correspond to anything we would ordinarily regard as a piece of reasoning. For example:

Snow is green.

_____

It has been a wet winter.

This generosity when it comes to counting things as arguments, while it might initially seem silly, is in fact good, for the following reason. As we shall discuss in the next section, one of our aims is to develop an account that will enable us to determine of any piece of reasoning—no matter what its subject matter—whether it is *valid*. (We shall see what validity is, and why it is important, below.) The more things we count as arguments, the more widely applicable our account of validity will be. If we were more stringent about what counts as an argument, then there would be a worry that some piece of reasoning to which we want our account to apply cannot be represented as an argument (in the more restricted sense) and so would be left out of account. Our present approach avoids this worry. All we are assuming is that any piece of reasoning can be represented as a sequence of propositions (an argument), one of which (the conclusion) is what the piece of reasoning is intended to establish, and the rest of which (the premises) are intended to provide support for that conclusion. That is, every piece of reasoning can be represented as an argument. The fact that the opposite does not hold—that not every argument (in our technical sense) corresponds to an ordinary piece of reasoning—will not matter.

1.3.1 Exercises

Represent the following lines of reasoning as arguments.

1. If the stock market crashes, thousands of experienced investors will lose a lot of money. So the stock market won't crash.

2. Diamond is harder than topaz, topaz is harder than quartz, quartz is harder than calcite, and calcite is harder than talc, therefore diamond is harder than talc.

3. Any friend of yours is a friend of mine; and you're friends with everyone on the volleyball team. Hence, if Sally's on the volleyball team, she's a friend of mine.

4. When a politician engages in shady business dealings, it ends up on page one of the newspapers. No South Australian senator has ever appeared on page one of a newspaper. Thus, no South Australian senator engages in shady business dealings.

1.4 Logical Consequence

Consider the following argument:

1. The rabbit ran down the left path or the right path.
   The rabbit did not run down the left path.
   ∴ The rabbit ran down the right path.

It is said that dogs exhibit a grasp of logic by reasoning in this way.[18] Suppose a dog is chasing a rabbit through the forest, when it comes to a fork in the path. The dog does not know which way the rabbit has gone, but it knows (because the undergrowth is impenetrable) that it has gone left or right (first premise). The dog sniffs down one path—say, the left one—trying to pick up the scent. If it does not pick up the scent, then it knows the rabbit has not gone down the left path (second premise). In this case the dog simply runs down the right path, without stopping to pick up the scent. For the dog knows, purely on the basis of logic—that is, without having to determine so by sniffing—that the rabbit has gone right: it must have, because it had to go left or right, and it did not go left, so that leaves only the possibility that it went right.

The argument is a good one. What exactly is good about it? Well, two things. The first is that given that the premises are true, there is no possibility of the conclusion's not being true. We can put the point in various ways:

- The truth of the premises guarantees the truth of the conclusion.

- It is impossible for the premises all to be true and the conclusion not be true.

- There is no way for the premises all to be true without the conclusion being true.

We call this property—the property that an argument has when it is impossible for the premises to be true and the conclusion false—*necessary truth-preservation* (NTP), and we call an argument with this property *necessarily truth-preserving* (NTP).[19]

Consider another example:[20]

2. All kelpies are dogs.
   Maisie is a dog.
   ∴ Maisie is a kelpie.

Figure 1.1. The argument is valid

Can we imagine a situation in which the premises are both true but the conclusion is false? Yes: suppose that (as in actual fact) all kelpies are dogs (so the first premise is true) and suppose that Maisie is a beagle (and hence a dog—so the second premise is true); in this case the conclusion is false. Hence argument (2) is not NTP.

Now consider a third example:

3. All kelpies are dogs.
   Maisie is a kelpie.
   ∴ Maisie is a dog.

Can we imagine a situation in which the premises are both true but the conclusion is false? No. Supposing the first premise to be true means supposing that (to represent the situation visually) a line drawn around all kelpies would never cross outside a line drawn around all dogs (Figure 1.1). Supposing the second premise to be true means supposing that Maisie is inside the line drawn around all kelpies. But then it is impossible for Maisie to be outside the line drawn around the dogs—that is, it is impossible for the conclusion to be false. So argument (3) is NTP.

There is a second good thing about argument (1), apart from its being NTP. Consider four more arguments:

4. Tangles is gray, and Maisie is furry.
   ∴ Maisie is furry.

5. Philosophy is interesting, and logic is rewarding.
   ∴ Logic is rewarding.

6. John is Susan's brother.
   ∴ Susan is John's sister.

7. The glass on the table contains water.
   ∴ The glass on the table contains H_2O.

All these arguments are NTP—but let's consider why each argument is NTP: what it is, in each case, that *underwrites* the fact that the premises cannot be true while the conclusion is false.

In the case of argument (4), it is the form or structure of the argument that makes it NTP.[21] The argument is a complex structure, built from propositions which themselves have parts. It is the particular way in which these parts are arranged to form the argument—that is, the form or structure of the argument—that ensures it is NTP. For the premise to be true, two things must be the case: that Tangles is gray, and that Maisie is furry. The conclusion claims that the second of these two things is the case: that Maisie is furry. Clearly, there is no way for the premise to be true without the conclusion being true. We can see this without knowing what Tangles and Maisie are (cats, dogs, hamsters—it doesn't matter). In fact, we do not even have to know what "gray" and "furry" mean. We can see that whatever Tangles and Maisie are and whatever properties "gray" and "furry" pick out, if it is true that Tangles is gray and Maisie is furry, then it must be true that Maisie is furry—for part of what it takes for "Tangles is gray, and Maisie is furry" to be true is that "Maisie is furry" is true.

The same can be said about argument (5). One does not have to know what philosophy and logic are—or what it takes for something to be interesting or rewarding—to see that if the premise is true, then the conclusion must be true, too: for part of what it takes for "philosophy is interesting and logic is rewarding" to be true is that "logic is rewarding" is true. Indeed it is clear that any argument will be valid if it has the following form, where A and B are propositions:

A and B
———————
B

It doesn't matter what propositions we put in for A and B: we could go through the same reasoning as above (the conclusion's being true is part of what it takes for the premise to be true) and thereby convince ourselves that the argument is valid.

Contrast arguments (6) and (7). In the case of (6), to see that the premise cannot be true while the conclusion is false, we need specific knowledge of the meanings of the terms involved. We need to know that "Susan" is a girl's name,[22] and that the meanings of "brother" and "sister" are related in a particular way: if a person x is the brother of a female y, then y is the sister of x. Accordingly, if we replace these terms with terms having different particular meanings, then the resulting arguments need not be NTP. For example:

8. John is Susan's friend.
   ∴ Susan is John's aunt.

9. John is Bill's brother.
   ∴ Bill is John's sister.

Contrast argument (4), where we could replace "Tangles" and "Maisie" with any other names, and "gray" and "furry" with terms for any other properties, and the resulting argument would still be NTP. For example:

10. Bill is boring, and Ben is asleep.
    ∴ Ben is asleep.

11. Jill is snoring, and Jack is awake.
    ∴ Jack is awake.

In the case of (7), to see that the premise cannot be true while the conclusion is false, we need specific scientific knowledge: we need to know that the chemical composition of water is H_2O.[23] Accordingly, if we replace the term "water" with a term for a substance with different chemical properties—or the term "H_2O" with a term for a different chemical compound—then the resulting arguments need not be NTP. For example:

12. The glass on the table contains sand.
    ∴ The glass on the table contains H_2O.

13. The glass on the table contains water.
    ∴ The glass on the table contains N_2O.

So, some arguments that are NTP are so by virtue of their form or structure: simply given the way the argument is constructed, there is no way for the premises to be true and the conclusion false. Other arguments that are NTP are not so by virtue of their form or structure: the way in which the argument is constructed does not guarantee that there is no way for the premises to be true and the conclusion false. Rather, the fact that there is no such way is underwritten by specific facts either about the meanings of the particular terms in the argument (e.g., "Susan"—this has to be a girl's name), or about the particular things in the world that these terms pick out (e.g., water—its chemical composition is H_2O), or both.

If an argument is NTP by virtue of its form or structure, then we call it *valid*, and we say that the conclusion is a *logical consequence* of the premises. There are therefore two aspects to validity/logical consequence:

1. The premises cannot be true while the conclusion is false (NTP).

2. The form or structure of the argument guarantees that it is NTP.

An argument that is not valid is said to be *invalid*. An argument might be invalid because it is not NTP, or because, although it is NTP, this fact is not underwritten by the structure of the argument.

Note that the foregoing does not constitute a precise definition of validity: it is a statement of a fundamental intuitive idea. One of our goals is to come up with a precise analysis of validity or logical consequence.[24] The guiding idea that we have set out—according to which validity is NTP by virtue of form—can be found, for example, in Alfred Tarski's seminal discussion of logical consequence, where it is presented as the traditional, intuitive conception:[25]

> I emphasize . . . that the proposed treatment of the concept of consequence makes no very high claim to complete originality. The ideas involved in this treatment will certainly seem to be something well known. . . . Certain considerations of an intuitive nature will form our starting-point. Consider any class K of sentences and a sentence X which follows from the sentences of this class. From an intuitive standpoint it can never happen that both the class K consists only of true sentences and the sentence X is false.[26] Moreover, since we are concerned here with the concept of logical, i.e., *formal*, consequence, and thus with a relation which is to be uniquely determined by the form of the sentences between which it holds, this relation cannot be influenced in any way by empirical knowledge, and in particular by knowledge of the objects to which the sentence X or the sentences of the class K refer.[27] . . . The two circumstances just indicated[28] . . . seem to be very characteristic and essential for the proper concept of consequence. [Tarski, 1936, 414–15]

Indeed, the idea goes back to Aristotle [c. 350 BC-b, §1], who begins by stating: "A deduction is a discourse in which, certain things being stated, something other than what is stated follows of necessity from their being so." This is the idea of NTP. Then, when discussing arguments, Aristotle first presents an argument form in an abstract way, with schematic letters in place of particular terms, for example:

> Every C is B.
> No B is A.
> Therefore no C is A.

He then derives specific arguments by putting particular terms in place of the letters, for example:

> Every swan is white.
> No white thing is a raven.
> Therefore no swan is a raven.

The reasoning that shows the argument to be NTP is carried out at the level of the argument form (i.e., in terms of As, Bs and Cs; not ravens, white things,

and swans): it is thus clear that Aristotle is interested in those arguments that are NTP by virtue of their form.

In this section, we have considered a number of examples of arguments and asked whether they are valid. We worked in an intuitive way, asking whether we could imagine situations in which the premises are true and the conclusion false. This approach is far from ideal. Suppose someone claims that she cannot imagine a situation in which the premises of argument (2) are true and the conclusion is false—or that someone claims that he can imagine a situation in which the premises of argument (3) are true and the conclusion is false. What are we to say in response? Can we show that these persons are mistaken? What we should like to have is a foolproof method of determining whether a given argument is valid: a method that establishes beyond doubt whether the argument is valid and that can be followed in a straightforward, routine way, without recourse to intuition or imagination. Think of the way you convince someone that $1{,}257 + 2{,}874 = 4{,}131$. You do not appeal to their imagination or intuition: you go through the mechanical process of long addition, first writing the numbers one above the other, then adding the units and carrying 1, then adding the tens, and so on, until the answer is attained. The task is thus broken up, in a specified way, into a sequence of small steps (adding numbers less than ten and carrying single digits), each of which is simple and routine. What we would like in the case of validity is something similar: a set of simple rules that we can apply in a specified order to a given argument, leading eventually to the correct verdict: valid or invalid.[29]

§

Recall the quotation from Peirce in §1.1 which ends "reasoning is good if it be such as to give a true conclusion from true premises, and not otherwise." The property of reasoning that Peirce mentions here—being such as to give a true conclusion from true premises—is NTP. In this passage, Peirce equates NTP with good reasoning. That view seems too strong—if "good reasoning" is taken to have its ordinary, intuitive meaning. For example, suppose that someone believes that there is water in the glass, but does not go on to conclude that there is H_2O in the glass. This does not necessarily mean that there is something wrong with her powers of reasoning: she may be fully rational, but simply not know that water is H_2O. Such a person could perhaps be criticized for not knowing basic science—but only if she could have been expected to know it (say, because she had completed high school)—but it would not be right to say that she had failed to reason well.

So we cannot equate good reasoning with NTP. Can we equate good reasoning with validity (i.e., NTP by virtue of form)? This suggestion might seem plausible at first sight. For example, if someone believes that Bill is boring and Ben is asleep, but he does not believe that Ben is asleep, then it seems that there

is certainly something wrong with his powers of reasoning. Yet even the claim that reasoning is good if and only if it is valid (as opposed to simply NTP) would be too strong. As we shall see in §1.5, an argument can be valid without being a good argument (intuitively speaking). Conversely, many good pieces of reasoning (intuitively speaking) are not valid, for the truth of the premises does not guarantee the truth of the conclusion: it only makes the conclusion highly probable.

Reasoning in which validity is a prerequisite for goodness is often called *deductive* reasoning. Important subclasses of nondeductive reasoning are *inductive* reasoning—where one draws conclusions about future events based on past observations (e.g., the sun has risen on every morning that I have experienced, therefore it will rise tomorrow morning), or draws general conclusions based on observations of specific instances (e.g., every lump of sugar that I have put in tea dissolves, therefore all sugar is soluble)—and *abductive* reasoning, also known as (aka) "inference to the best explanation"—where one reasons from the data at hand to the best available explanation of that data (e.g., concluding that the butler did it, because this hypothesis best fits the available clues).[30] Whereas validity is a criterion of goodness for deductive arguments, the analogous criterion of goodness for nondeductive arguments is often called *inductive strength:* an argument is inductively strong just in case it is *improbable*—as opposed to *impossible,* in the case of validity—that its premises be true and its conclusion false.

The full story of the relationship between validity and good reasoning is evidently rather complex. It is not a story we shall try to tell here, for our topic is logic—and as we have noted, logic is the science of truth, not the science of reasoning. However, this much certainly seems true: if we are interested in reasoning—and in classifying it as good or bad—then one question of interest will always be "is the reasoning valid?" This is true regardless of whether we are considering deductive or nondeductive reasoning. The answer to the question "is the reasoning valid?" will not, in general, completely close the issue of whether the reasoning is good—but it will never be irrelevant to that issue. Therefore, if we are to apply logic—the laws of truth—to the study of reasoning, it will be useful to be able to determine of any argument—no matter what its subject matter—whether it is valid.

§

When it comes to validity, then, we now have two goals on the table. One is to find a precise analysis of validity. (Thus far we have given only a rough, guiding idea of what validity is: NTP guaranteed by form. As we noted, this does not amount to a precise analysis.) The other is to find a method of assessing arguments for validity that is both

1. foolproof: it can be followed in a straightforward, routine way, without recourse to intuition or imagination—and it always gives the right answer; and

2. general: it can be applied to any argument.

Note that there will be an intimate connection between the role of form in the definition of validity (an argument is valid if it is NTP by virtue of its form) and the goal of finding a method of assessing arguments for validity that can be applied to any argument, no matter what its subject matter. It is the fact that validity can be assessed on the basis of form, in abstraction from the specific content of the propositions involved in an argument (i.e., the specific claims made about the world—what ways, exactly, the propositions that make up the argument are representing the world to be), that will bring this goal within reach.

1.4.1 Exercises

State whether each of the following arguments is valid or invalid.

1. All dogs are mammals.
   All mammals are animals.
   _____

   All dogs are animals.

2. All dogs are mammals.
   All dogs are animals.
   _____

   All mammals are animals.

3. All dogs are mammals.
   No fish are mammals.
   _____

   No fish are dogs.

4. All fish are mammals.
   All mammals are robots.
   _____

   All fish are robots.

1.5 Soundness

Consider argument (4) in Exercises 1.4.1. It is valid, but there is still something wrong with it: it does not establish its conclusion as true—because its premises are not in fact true. It has the property that if its premises were both true, then its conclusion would have to be true—that is, it is NTP—but its premises are not in fact true, and so the argument does not establish the truth of its conclusion.

We say that an argument is *sound* if it is valid and, in addition, has premises that are all in fact true:

sound = valid + all premises true

A valid argument can have any combination of true and false premises and conclusion except true premises and a false conclusion. A sound argument has true premises and therefore—because it is valid—a true conclusion.

Logic has very little to say about soundness—because it has very little to say about the actual truth or falsity of particular propositions. Logic, as we have said, is concerned with the laws of truth—and the general laws of truth are very different from the mass of particular facts of truth, that is, the facts as to which propositions actually are true and which are false. There are countless propositions concerning all manner of different things: "two plus two equals four," "the mother of the driver of the bus I caught this morning was born in Cygnet," "the number of polar bears born in 1942 was 9,125," and so on. No science would hope to tell us whether every single one is true or false. This is not simply because there are too many of them: it is in the nature of science not to catalogue particular matters of fact but to look for interesting patterns and generalizations—for laws. Consider physics, which is concerned (in part) with motion. Physicists look for the general laws governing all motions: they do not seek to determine all the particular facts concerning what moves where, when, and at what speeds. Of course, given the general laws of motion and some particular facts (e.g., concerning the moon's orbit, the launch trajectory of a certain rocket, and a number of other facts) one can deduce other facts (e.g., that the rocket will reach the moon at such and such a time). The same thing happens in logic. Given the general laws of truth and some particular facts (e.g., that this proposition is true and that one is false) one can deduce other facts (e.g., that a third proposition is true). But just as it is not the job of the physicist to tell us where everything is at every moment and how fast it is moving, so too it is not the job of the logician to tell us whether every proposition is true or false. Therefore, questions of soundness—which require, for their answer, knowledge of whether certain premises are actually true—fall outside the scope of logic.[31]

Likewise, logic is not concerned with whether we know that the premises of an argument are true. We might have an argument that includes the premise "the highest speed reached by a polar bear on 11 January 2004 was 31.35 kilometres per hour." The argument might (as it happens) be sound, but that would not make it a convincing argument for its conclusion, because we could never know that all the premises were true.

So it takes more than validity to make a piece of (deductive) reasoning convincing. A really convincing argument will be not only valid but also sound,

and furthermore have premises that can be known to be true. Some people have complained that logic—which tells us only about validity—does not provide a complete account of good reasoning. It is entirely true that logic does not provide a complete account of good reasoning, but this is cause for complaint only if one thinks that logic is the science of reasoning. From our perspective there is no problem here: logic is not the science of reasoning; it is the science of truth. Logic has important applications to reasoning—most notably in what it says about validity. However there is both more and less to good reasoning than validity (i.e., valid arguments are not always good, and good arguments are not always valid)—and hence (as already noted in §1.4) there is more to be said about reasoning than can be deduced from the laws of truth.

1.5.1 Exercises

1. Which of the arguments in Exercises 1.4.1 are sound?

2. Find an argument in Exercises 1.4.1 that has all true premises and a true conclusion but is not valid and hence not sound.

3. Find an argument in Exercises 1.4.1 that has false premises and a false conclusion but is valid.

1.6 Connectives

We have said that an argument is valid if its structure guarantees that it is NTP. It follows immediately that if validity is to be an interesting and useful concept, some propositions must have internal structure. For suppose that all propositions were simply "dots," with no structure. Then the only valid arguments would be ones where the conclusion is one of the premises. That would render the concept of validity virtually useless and deprive it of all interest. We are going to assume, then, that at least some propositions have internal structure—and of course this assumption is extremely natural. Consider our argument:

> Tangles is gray, and Maisie is furry.
> ∴ Maisie is furry.

It seems obvious that the first premise—so far from being a featureless dot with no internal structure—is a proposition made up (in some way to be investigated) from two other propositions: "Tangles is gray" and "Maisie is furry."

Before we can say anything useful about the forms that arguments may take, then, our first step must be to investigate the internal structure of the things

that make up arguments—that is, of propositions. We divide propositions into two kinds:

1. basic propositions: propositions having no parts that are themselves propositions; and

2. compound propositions: propositions made up from other propositions and *connectives.*

In Part I of this book—which concerns propositional logic—we look at the internal structure of compound propositions, that is, at the ways in which propositions may be combined with connectives to form larger propositions. It will not be until Part II—which concerns predicate logic—that we shall look at the internal structure of basic propositions.

A compound proposition is made up of component propositions and connectives. We now embark upon an investigation of connectives. Our investigation will be guided by our interest in the laws of truth. We saw that any argument of the form "A and B /∴ B" is valid (§1.4). The reason is that the premise is made up of two component propositions, A and B, put together by means of "and"—that is, in such a way that the compound proposition can be true only if both components are true—and the conclusion is one of those component propositions. Hence, the conclusion's being true is part of what it takes for the premise to be true. Thus, if the premise is true, the conclusion must be too: the laws of truth ensure, so to speak, that truth flows from premise to conclusion (if it is present in the premise in the first place). So the validity of the argument turns on the internal structure of the premise—in particular, on the way that the connective "and" works in relationship to truth and falsity.

Our search for connectives will be guided by the idea that we are interested only in those aspects of the internal structure of compound propositions that have an important relationship to truth and falsity. More specifically, we shall focus on a particular kind of relationship to truth and falsity: the kind where the truth or falsity of the compound proposition depends solely on the truth or falsity of its component propositions. A connective is *truth functional* if it has the property that the truth or falsity of a compound proposition formed from the connective and some other propositions is completely determined by the truth or falsity of those component propositions. Our focus, then, will be on truth-functional connectives.[32]

1.6.1 Negation

Consider the proposition "Maisie is not a rottweiler." Thinking in terms of truth and falsity, we can see this proposition as being made up of a component proposition ("Maisie is a rottweiler") and a connective (expressed by "not")

that have the following relationship to one another: if "Maisie is a rottweiler" is true, then "Maisie is not a rottweiler" is false, and if "Maisie is a rottweiler" is false, then "Maisie is not a rottweiler" is true. We use the term "negation" for the connective which has this property: viz. it goes together with a proposition to make up a compound proposition that is true just in case the component proposition is false. Here is some terminology:

"Maisie is not a rottweiler" is the *negation* of "Maisie is a rottweiler."
"Maisie is a rottweiler" is the *negand* of "Maisie is not a rottweiler."

Note the double meaning of negation. On the one hand we use it to refer to the connective, which goes together with a proposition to make up a compound proposition. On the other hand we use it to refer to that compound proposition. This ambiguity is perhaps unfortunate, but it is so well entrenched in the literature that we shall not try to introduce new terms here. As long as we are on the lookout for this ambiguity, it should not cause us any problems.

Using our new terminology, we can express the key relationship between negation and truth this way:

If the negand is true, the negation is false, and if the negand is false, the negation is true.

Thus, negation is a truth-functional connective: to know whether a negation is true or false you need only know whether the negand is true or false: the truth or falsity of the negation is completely determined by the truth or falsity of the negand.

It is this particular relationship between negation and truth—rather than the presence of the word "not"—that is the defining feature of negation. Negation can also be expressed in other ways, for example:

- It is not the case that there is an elephant in the room.

- There is no elephant in the room.

- There is not an elephant in the room.

All these examples can be regarded as expressing the negation of "there is an elephant in the room."

Connectives can be applied to any proposition, basic or compound. Thus, we can negate "Bob is a good student" to get "Bob is not a good student," and we can also negate the latter to get "it is not the case that Bob is not a good student," which is sometimes referred to as the *double negation* of "Bob is a good student."

To get a complete proposition using the connective negation, we need to add the connective to one proposition (the negand). Thus negation is a *one-place* (aka *unary* or *monadic*) connective.

1.6.1.1 EXERCISES

1. What is the negand of:
     (i) Bob is not a good student
    (ii) I haven't decided not to go to the party.
   (iii) Mars isn't the closest planet to the sun.
    (iv) It is not the case that Alice is late.
     (v) I don't like scrambled eggs.
    (vi) Scrambled eggs aren't good for you.

2. If a proposition is true, its double negation is . . . ?

3. If a proposition's double negation is false, the proposition is . . . ?

1.6.2 Conjunction

Consider the proposition "Maisie is tired, and the road is long." Thinking in terms of truth and falsity, we can see this proposition as being made up of two component propositions ("Maisie is tired" and "the road is long") and a connective (expressed by "and"), which have the following relationship to one another: "Maisie is tired, and the road is long" is true just in case "Maisie is tired" and "the road is long" are both true. We use the term *conjunction* for the connective that has this property: it goes together with two propositions to make up a compound proposition that is true just in case both component propositions are true.[33] Here is some terminology:

> "Maisie is tired and the road is long" is the *conjunction* of "Maisie is tired" and
> "the road is long."
> "Maisie is tired" and "the road is long" are the *conjuncts* of "Maisie is tired
> and the road is long."

Again, "conjunction" is used in two senses: to pick out a connective and to pick out a compound proposition built up using this connective.

Using our new terminology, we can express the key relationship between conjunction and truth in this way:

> The conjunction is true just in case both conjuncts are true.
> If one or more of the conjuncts is false, the conjunction is false.

Thus, conjunction is a truth-functional connective: to know whether a conjunction is true or false you need only know whether the conjuncts are true or

false: the truth or falsity of the conjunction is completely determined by the truth or falsity of the conjuncts.

It is this particular relationship between conjunction and truth—rather than the presence of the word "and"—that is the defining feature of conjunction. Conjunction can also be expressed in other ways, and not every use of "and" in English expresses truth-functional conjunction. For the moment, however, we shall continue with our preliminary examination of truth-functional connectives; we turn to a detailed discussion of the relationships among these connectives and expressions of English in Chapter 6. So keep in mind throughout the remainder of this chapter: we are here giving a first, brief introduction to truth-functional connectives via English words that typically, often, or sometimes express these connectives. In Chapters 2 and 3 we shall gain a much deeper understanding of these connectives via the study of a new symbolic language before returning to the subtleties of the relationships between these connectives and expressions of English in Chapter 6.

To obtain a complete proposition using the conjunction connective, we need to add the connective to two propositions (the conjuncts). Thus, conjunction is called a *two-place* (aka *binary* or *dyadic*) connective.

1.6.2.1 EXERCISES
What are the conjuncts of the following propositions?

1. The sun is shining, and I am happy.

2. Maisie and Rosie are my friends.

3. Sailing is fun, and snowboarding is too.

4. We watched the movie and ate popcorn.

5. Sue does not want the red bicycle, and she does not like the blue one.

6. The road to the campsite is long and uneven.

1.6.3 Disjunction
Consider the proposition "Frances had eggs for breakfast or for lunch." Thinking in terms of truth and falsity, we can see this proposition as being made up of two component propositions ("Frances had eggs for breakfast" and "Frances had eggs for lunch") and a connective (expressed by "or"), which have the following relationship to one another: "Frances had eggs for breakfast or for lunch" is true just in case at least one of "Frances had eggs for breakfast" and "Frances had eggs for lunch" are true. We use the term *disjunction* for the connective that has this property: it goes together with two propositions to

make up a compound proposition that is true just in case at least one of those component propositions is true. Here is some terminology:

> "Frances had eggs for breakfast or for lunch" is the *disjunction* of "Frances had eggs for breakfast" and "Frances had eggs for lunch."
> "Frances had eggs for breakfast" and "Frances had eggs for lunch" are the *disjuncts* of "Frances had eggs for breakfast or for lunch."

Using this terminology, we can express the key relationship between disjunction and truth in this way:

> The disjunction is true just in case at least one of the disjuncts is true.
> If both the disjuncts are false, the disjunction is false.

Thus, disjunction is a truth-functional connective: to know whether a disjunction is true or false you need only know whether the disjuncts are true or false: the truth or falsity of the disjunction is completely determined by the truth or falsity of the disjuncts.

It is this relationship between disjunction and truth—rather than the use of the word "or" as in the example above—that is the defining feature of disjunction. Disjunction can also be expressed in other ways, for example:

- Either Frances had eggs for breakfast or she had eggs for lunch.

- Frances had eggs for breakfast and/or lunch.

- Frances had eggs for breakfast or lunch—or both.

To obtain a complete proposition using the disjunction connective, we need to add the connective to two propositions (the disjuncts). Thus, disjunction is a two-place connective.

1.6.4 Conditional

Imagine that we look out the window and see a haze; we are not sure whether it is smoke, fog, dust, or something else. Consider the proposition "if that is smoke, then there is a fire." A proposition of this form has two components, and claims that if one of them is true, then the other is true too. We call the former component the *antecedent,* the latter component the *consequent,* and the compound proposition a *conditional.* (Once again we also use the term "conditional" for the two-place connective used to form this compound proposition.) In the above example the conditional is "if that is smoke, then there is a fire," the antecedent is "that is smoke," and the consequent is "there is a fire."

Note that the antecedent is not always written first. The antecedent is the component proposition of which it is said that if it is true, then another proposition is true; the consequent is that other proposition. To put it another way: if the conditional is true, then one of its components might be true without the other component being true, but not vice versa. The consequent is the component that might be true even if the other component is not true; the antecedent is the component that cannot be true without the other component also being true (assuming the conditional as a whole is true). Thus, the relationship between the antecedent and the consequent is logical or *alethic* (having to do with truth), not temporal or spatial. If I say "there is a fire if that is smoke," the antecedent is "that is smoke," and the consequent is "there is a fire." In other words, this is just a different way of expressing the same conditional.

As well as being expressed by "if . . . then" and "if," conditionals can also be expressed using "only if." For example, suppose that I have just gotten off a mystery flight and am wondering where I am. Consider the proposition "I am in New York only if I am in America." This is a conditional in which the antecedent is "I am in New York," and the consequent is "I am in America:" it thus says the same thing as "if I am in New York, I am in America." The easiest way to see this is to think what it would take to make the latter claim false: I would have to be in New York without being in America. So "if I am in New York, I am in America" rules out the case in which I am in New York but am not in America. And that is exactly what "I am in New York only if I am in America" does: the claim is that it does not happen that I am in New York but not in America. In contrast, "I am in America only if I am in New York" says something quite different: it says the same thing as "if I am in America, then I am in New York." In general, "if P then Q" and "P only if Q" say the same thing.

1.6.4.1 EXERCISES

What are the (a) antecedents and (b) consequents of the following propositions?

1. If that's pistachio ice cream, it doesn't taste the way it should.

2. That tastes the way it should only if it isn't pistachio ice cream.

3. If that is supposed to taste that way, then it isn't pistachio ice cream.

4. If you pressed the red button, then your cup contains coffee.

5. Your cup does not contain coffee if you pressed the green button.

6. Your cup contains hot chocolate only if you pressed the green button.

1.6.5 Biconditional

Suppose the drink machine has unlabeled buttons, and you are wondering what is in your cup, which you have just removed from the machine. Consider the proposition "your cup contains coffee if and only if you pressed the red button." Someone who asserts this is committed to two claims:

Your cup contains coffee if you pressed the red button.
Your cup contains coffee only if you pressed the red button.

The first is a conditional with antecedent "you pressed the red button" and consequent "your cup contains coffee." The second is a conditional with antecedent "your cup contains coffee" and consequent "you pressed the red button." Now, under what conditions is the original proposition true? Suppose your cup contains coffee. Then, if the second conditional is to be true, it must be the case that you pressed the red button. Suppose your cup does not contain coffee. Then, if the first conditional is to be true, it must be the case that you did not press the red button. So the original proposition ("your cup contains coffee if and only if you pressed the red button") is true if your cup contains coffee and you pressed the red button, and true if your cup does not contain coffee and you did not press the red button, but it is false if your cup contains coffee and you did not press the red button, and it is false if your cup does not contain coffee and you did press the red button. In other words, it is true just in case its two component propositions ("your cup contains coffee" and "you pressed the red button") have the same truth value—that is, are both true, or both false.

We call the original claim a *biconditional.* Note that we are here regarding the proposition "your cup contains coffee if and only if you pressed the red button" as formed from two propositions ("your cup contains coffee" and "you pressed the red button") using the two-place connective "if and only if," that is, the biconditional. We regard this claim as equivalent to the conjunction of the two conditionals "your cup contains coffee if you pressed the red button" and "your cup contains coffee only if you pressed the red button"—but it is not the same proposition as "your cup contains coffee if you pressed the red button and your cup contains coffee only if you pressed the red button." The latter is a compound proposition built up using two basic propositions ("your cup contains coffee" and "you pressed the red button") and two different connectives (a conditional used twice and a conjunction). This idea of different propositions being equivalent—that is, true and false in the same situations—will be made clear in §4.3.[34]

Note that it is common to abbreviate "if and only if" as "iff," and that "just in case" is often used as a synonym for "if and only if" (e.g., "a conjunction is true

just in case both its conjuncts are true" states the same thing as "a conjunction is true if and only if both its conjuncts are true").

1.6.6 Exercises

State what sort of compound proposition each of the following is, and identify its components. Do the same for the components.

1. If it is sunny and windy tomorrow, we shall go sailing or kite flying.

2. If it rains or snows tomorrow, we shall not go sailing or kite flying.

3. Either he'll stay here and we'll come back and collect him later, or he'll come with us and we'll all come back together.

4. Jane is a talented painter and a wonderful sculptor, and if she remains interested in art, her work will one day be of the highest quality.

5. It's not the case that the unemployment rate will both increase and decrease in the next quarter.

6. Your sunburn will get worse and become painful if you don't stop swimming during the daytime.

7. Either Steven won't get the job, or I'll leave and all my clients will leave.

8. The Tigers will not lose if and only if both Thompson and Thomson get injured.

9. Fido will wag his tail if you give him dinner at 6 this evening, and if you don't, then he will bark.

10. It will rain or snow today—or else it won't.

INDEX