

1 * Introduction

1. An Example

This chapter introduces the subject of formal logic. Toward this end it seems natural to begin with a definition, one that explains what formal logic is. This is the normal practice in teaching a whole range of specialized fields, but in philosophy a definition is usually a false start, at least it is if we take the definition seriously and plan to stick to it. A definition determines how the thing being defined is to be understood. In philosophy, however, we must first come to some preliminary understanding of the thing (in some broad sense of the word ‘thing’), an understanding that is then critically explored and perhaps altered. For this purpose, the device of the initial definition, held constant through all subsequent discussion, is ill suited.

In philosophy, we generally start out with some prior knowledge of the subject under study, albeit not the sort of knowledge we really want. Frequently, we find our subject somehow implicit in some more-or-less concrete fact or example. The analysis of such examples is thus an extremely useful way of approaching the subject itself. In this way we shall approach one of the central notions of formal logic, the notion of valid inference. As to why this notion is central to formal logic, and why this logic is called “formal,” these are questions we cannot yet answer. But our analysis of examples will get us started in the right direction.

I begin with an example of the familiar, trivial variety so often employed in logic:

Example 1.1 All logicians are human
 All humans need sleep
 *
 Therefore, All logicians need sleep

To simplify our discussion of this and future examples, we will call the sentences above the asterisk “premises” and the sentence below it the “conclusion.” Being a premise or a conclusion is clearly not a property of a sentence in and of itself—it is merely a function of whether a given sentence is found above or below the asterisk. Now, in this case the move from premises to conclusion is somehow compelling, and so we call the whole a (logically) valid inference. It is of the utmost importance to realize that the expression “valid inference” refers to the entire system of three sentences and not merely to the conclusion. In other words, validity is a feature of the relationship between premises and conclusion and is not a property of the conclusion itself. Of course there are things that *can* be said about the conclusion, such as that it is true, but then we are not talking about the inference, which again is the particular relationship between premises and conclusion. I will return to this distinction in greater detail below (see II.2.4.b), but for now suffice it to insist that the following are two separate questions:

1. Is the inference valid?
2. Is the conclusion true?

(The ease with which we confuse these questions is due to the fact that there is, after all, a connection between the validity of an inference and the truth of its conclusion; I will return to this later.) As we shall see, the distinction between these two questions is central to logic.

It is a truly remarkable fact that one rarely encounters anyone who fails to agree that Example 1.1 constitutes a logically valid inference, in the sense that the transition from premises to conclusion is somehow compelling. Without having to come to any agreement, or demand any sort of explanation, anyone with sufficient mastery of the English language to un-

derstand the sentences in question recognizes the inference as valid. We would be hard pressed to find someone who, having accepted the premises, would refuse to accept the conclusion. It is worth asking why this is so. At this stage, however, I cannot pursue this question. Let us extract five features of logically valid inferences from our example. Anyone who accepts Example 1.1 as logically valid will also be prepared to ascribe all of these features to it and to related inferences.

Feature 1. When the premises are true, then the conclusion is also true.

Here we must pay close attention to what is being said. We are not claiming *that the premises are true*, but only that *if they are true, then the conclusion is also true*. This feature holds with all logically valid inferences and thus deserves a name. I shall call it the “truth-transferring” property of logically valid inferences. So, when the premises of a logically valid inference happen to be true, the conclusion is also true—the truth of the premises is *transferred* to the conclusion.¹

This truth-transferring quality of our example is, in the end, something one must see for oneself; anyone incapable of seeing it will never understand what logic is all about. Thus, as far as this fundamental insight into truth transfer is concerned, logic cannot be taught. To be sure, I can hint and gesture toward the insight, toward the experience of a certain thought process, with some degree of clarity. But no such descriptions of the relevant thought process will serve as a set of instructions such that, if one only follows the recipe, one will come to understand what truth transfer is. Experience of the thought process in question cannot be replaced by descriptions of it, and this experience is fundamental to logic. In this

1. Here the limitations of the truth transfer metaphor become apparent. When funds are transferred from account *A* to account *B*, they are obviously no longer in account *A*. By contrast, successful truth transfer from premises to conclusion does not mean that the premises are no longer true (provided they were true in the first place). Furthermore, in our present example, the conclusion “All logicians need sleep” is true anyway. No “transfer” of truth from the premises is required in order to make it true.

regard, our situation is far from unique. There are other experiences for which no description can be substituted, such as the experience of being in love, of the effects of alcohol, or of swimming in a strong current. Here, too, only someone who has had the right sort of experience will understand what the discussion is about.

Feature 2. The truth of the premises plays no role in our assessment of the validity of the inference.

Here it is claimed that the validity of an inference cannot be judged on the basis of the truth of its premises. Before defending this assertion, which may strike some as implausible, I hasten to point out that this second feature is not inconsistent with the first. Feature 1 is a claim about what happens *when* the premises are true; it does not assert *that* the premises are true. Now for the defense of Feature 2: Let us substitute the expression ‘have characteristic *S*’ for the expression ‘need sleep’ wherever the latter occurs in Example 1.1. This yields:

Example 1.2. All logicians are human.
 All humans have characteristic *S*.
 *
 Therefore, All logicians have characteristic *S*.

As is plain to see, we have another logically valid inference before us. But it must also be noted that the sentences ‘All humans have characteristic *S*’ and ‘All logicians have characteristic *S*’ are not straightforwardly true or false statements. For the meaning of the expression ‘characteristic *S*’ is indeterminate, thus leaving the two sentences in which it occurs equally indeterminate. Nonetheless, the logical validity of Example 1.2 may be seen without the slightest difficulty, even more easily than the validity of Example 1.1. Example 1.2 thus demonstrates that our assessment of the validity of an inference cannot depend on the truth of its premises, for in this example the second premise is so indeterminate as to be neither true nor false. Accordingly, ascertaining the validity of Example 1.1 does not re-

quire that we consult the Institute for Sleep Study on the distribution of sleep requirements among various segments of the population. The truth of the premises of this earlier example is equally irrelevant to its validity.

Indeed, the premises might be manifestly false without undermining the validity of the inference. This can easily be demonstrated by substituting the expression ‘have characteristic S ’ as it occurs in Example 1.2 with some other expression that makes the second premise false (for example, ‘are reptiles’). Now, one might be tempted to suppose that it is impossible to validly infer anything from false premises, but this temptation must be resisted. Arguments of the following form are, after all, commonplace: “Let us assume that all humans are reptiles. Under this assumption, what would logicians be?” In mathematics, inferences drawn from false premises are systematically employed in a procedure known as indirect proof, in which we assume the exact opposite of what we are trying to prove and infer a contradiction from it. This contradictory consequence shows that our assumption is false and thus that its opposite, the claim we were out to prove in the first place, must be true (for a more detailed exposition of indirect proof, see section II.2.4.f).

Feature 3. From any valid inference, many additional valid inferences may be generated (mechanically).

The procedure by which these additional inferences are generated has actually already put in an appearance, in our explication of Feature 2. Simply replace such expressions as ‘logician’, ‘human,’ and ‘need sleep’ as they occur in Example 1.1 with other expressions, such as ‘animal’, ‘life form’, and ‘God’s creation’, respectively. Our replacements must be methodical, in the sense that for *every* occurrence of an expression being replaced, we replace it with *one and the same* new expression; otherwise, our procedure will fail to generate further valid inferences. The substitution proposed above thus yields:

Example 1.3. All animals are life forms.
 All life forms are God's creations.
 *
 Therefore, All animals are God's creations.

This inference, too, is logically valid, as should be immediately clear. In this instance, the procedure of replacing each and every occurrence of a given expression with an occurrence of one and the same new expression has thus yielded a new valid inference. But of course this success hardly justifies the claim made in Feature 3, since it remains conceivable that the same procedure might fail for other examples. However, consideration of the fourth feature should persuade us that, in fact, our procedure will *always* yield new valid inferences.

Feature 4. The meanings of the expressions occurring in an inference are irrelevant to its validity.

The above formulation of Feature 4 is provisional and somewhat imprecise. By way of clarification, we return to Example 1.1, this time replacing every occurrence of 'logician', 'human', and 'need sleep' by *A*, *B*, and *C*, respectively, yielding:

Example 1.4. All *As* are *Bs*.
 All *Bs* are *Cs*.
 *
 Therefore, All *As* are *Cs*.

Once again, as with Example 1.2, we observe the logical validity of this inference, this despite the fact that *A*, *B*, and *C* are completely indeterminate. It follows that logically valid inference has nothing to do with word meaning, for the validity of an inference can be detected even after certain expressions ('logician', 'human', etc.) have been replaced with single letters. The "mechanism" of a valid inference thus depends not on the particular expressions it contains, but on something else. We will return to this something else later.

Having performed the abstraction procedure involved in converting Example 1.1 into Example 1.4, we begin to understand why Features 1–3 really *are* features of valid inferences. Feature 1 asserts that any valid inference is such that if its premises are true, then its conclusion is also true. In establishing the validity of Example 1.4, we consider this: if all *As* are *Bs*, and if all *Bs* are also *Cs*, is it then the case that all *As* are *Cs*? This question in essence asks whether the inference is truth-transferring—whether the truth of the premises is transferred to the conclusion. We can answer it in the affirmative without having the slightest idea what it would mean for all *As* to be *Bs* or all *Bs* to be *Cs*, let alone whether both claims—both premises—are actually true. But this is precisely what Feature 2 asserts of valid inferences. This result further underscores the possibility of using any valid inference to generate further valid inferences (Feature 3), for since the meanings of constituent expressions are irrelevant to the validity of the inference, they may be replaced with new expressions, provided we always replace each original expression with one and the same substitute. This reasoning establishes that Features 1–3 are implicit in Feature 4. Anyone who follows this reasoning is well on his or her way to grasping what logic is all about.

It should be noted that in our various manipulations of Example 1.1 we have played around with several of its constituent expressions, but not with all of them. For example, we have always left the word ‘all’ unchanged. Now let us see what happens when we replace the word ‘all’, as it occurs in Example 1.4, with the word ‘some’. This substitution yields:

Example 1.5. Some *As* are *Bs*
 Some *Bs* are *Cs*
 *
 Therefore, Some *As* are *Cs*

Is this inference valid, in the sense that the truth of the premises guarantees the truth of the conclusion? To simplify, let us replace ‘*A*’, ‘*B*’, and ‘*C*’ by particular meaningful expressions. If Example 1.5 is valid, then by

Feature 4, the result of this replacement should also be a valid inference. Let us substitute ‘plant’, ‘meat-eater’, and ‘cat’ for ‘ A ’, ‘ B ’, and ‘ C ’ respectively. This yields:

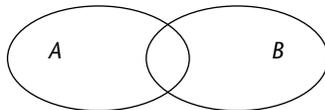
Example 1.6. Some plants are meat-eaters.
 Some meat-eaters are cats.
 *
 Therefore, Some plants are cats.

This conclusion is obviously false. But since both premises are true, their truth cannot have been transferred to the conclusion. The inference is thus not valid. Any inference that fails the test of validity is called “invalid.”

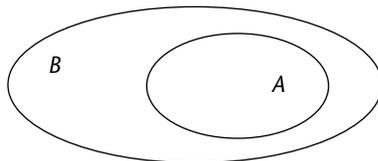
***Exercise 1.1.** Find other examples, analogous to 1.6, that demonstrate the invalidity of 1.5.

***Exercise 1.2.** Find expressions which, when inserted into Example 1.5, yield true premises *and* a true conclusion. How can such examples be reconciled with the invalidity of the inference?

***Exercise 1.3.** Let the claim “Some A s are B s” be represented by the following diagram:



Now let “All A s are B s” be represented by the following diagram:



Using these diagrams, explain why Example 1.4 constitutes a valid inference while 1.5 is invalid. Then analyze your solution to Exercise 1.2 by means of similar diagrams.

It has thus emerged that in a valid inference, such as Example 1.1, expressions like ‘human’, ‘need sleep’, and the like play a different role from expressions like ‘all’ (or ‘some’). Expressions of the first type can be replaced by other such expressions without affecting the validity of the inference (provided each occurrence of a given expression is always replaced by an occurrence of the same replacement). On the other hand, expressions of the second type cannot be so replaced; the validity of the inference depends on them. I will call this observation the fifth feature of valid inferences, noting that our formulation of Feature 5, like that of Feature 4, will have to be provisional:

Feature 5. The validity or invalidity of inferences depends on such expressions as ‘all’ and ‘some’.

**Exercise 1.4.* Formulate a valid inference containing the expression ‘some’ and three other expressions (or letters).

2. Preliminary Remarks on the Notion of Logical Form

The notion of logical form, to which we now turn, is of central importance to logic. Before going any further, I must first explain what we mean to achieve by introducing it. As we have just seen, the validity of an inference does not depend on every constituent expression in its component statements. On the contrary, the meanings of many such expressions, such as ‘logician’, ‘human’, and ‘need sleep’ in Example 1.1, are irrelevant to its validity. Any component of a statement that *does* contribute to the validity of inferences containing that statement is said to belong to the *logical form of the statement*. By contrast, any component of a statement that remains irrelevant to the validity of inferences containing the statement will be consigned to the (*logical*) *content of the statement*. For a preliminary explication of this distinction, we turn once again to Example 1.1. As we have seen, in the first premise, “All logicians are human,” while the meanings of ‘logician’ and ‘human’ were irrelevant to the subsequent in-

ference, the same was not true for ‘all’. The logical form of this statement might thus be rendered as ‘All *As* are *Bs*’. Similarly, ‘Some *As* are *Bs*’ is a good candidate for the logical form of ‘Some textbooks are exciting’. It thus appears that the move from Example 1.1 to Example 1.4 is just a matter of focusing in on logical forms for the constituent statements, where everything irrelevant to the validity of the inference has been stripped away from 1.4.

Our aim in introducing the notion of logical form will thus be to allow us to isolate exactly those components of statements that are relevant to the validity of inferences in which the statements occur. Success here will allow us to assess the validity of a given inference by reference only to the logical forms of its constituent statements, without having to engage with their contents. I hasten to point out that up to now I have merely articulated the main reason for wanting a notion of logical form, without doing more than gesturing toward the notion itself. It will not be properly introduced until sections II.1.2, and II.1.2. For the present, I content myself with four remarks on logical form.

1. The logical form of a sentence is obtained by *abstracting from the content* of the sentence. This is a straightforward consequence of our preliminary observations. “Abstraction” comes from the Latin *abstrahere*, meaning “to pull away” or “remove.” What is removed from a sentence is its content, and what is left is its logical form. This process draws a sharp line through a sentence. On one side of this line are the parts belonging to logical form; those on the other side belong to content. Graphic as this description is, by itself it remains indeterminate, for we need to know *precisely where to draw this line*. To say that we draw the line between form and content is plainly insufficient. Why, one might rejoin, should the ‘all’ in “All humans need sleep” belong to logical form rather than content? There are other sorts of lines one can draw between the forms and contents of sentences, as, for example, in grammar. There are several ways to abstract a grammatical form from the sentence “All humans need sleep.” Abstracting by syntactic components, we have modifier (‘all’), subject

(‘humans’), predicate (‘need’), direct object (‘sleep’). Or we might focus instead on the parts of speech, observing that the sentence is composed of an indefinite pronoun (‘all’), two nouns (‘humans’ and ‘sleep’), and a verb (‘need’). If we prefer, we can entirely abstract away word meanings to consider only the length of the sentence, expressed in the number of words. Finally, we might abstract by syntactic structure, in which case we would find that the sentence consists of a single independent clause.

As these examples plainly illustrate, there are many ways to turn our sentence into an abstraction. In each case, what is being *abstracted away* may be called the content of the sentence, and what is left over may be called its form. In order to target a *particular* form/content distinction, we must specify *precisely where the line between form and content should be drawn*. This specification will depend, in turn, on the exact nature of the task our form/content distinction is designed to perform. In ignorance of this task, the call for a form/content distinction remains vacuous. Proper attention to our reason for introducing the notion of logical form in the first place is thus indispensable to any real understanding of formal logic: motivated by our desire to analyze the logical validity of inferences, the notion of logical form allows us to isolate those aspects of sentences that are relevant to the validity of inferences in which those sentences occur. The somewhat less common designation of “(logical) content” applies to everything else.

2. But even when we confine ourselves to the logical form of a sentence, there is still more than one way to draw the line between form and content. The basis for this ambiguity will be treated with greater care later. For now, suffice it to say that any reference to *the* logical form of a sentence must be viewed with suspicion, for every sentence has more than one logical form. There are different subfields of logic, each of which is characterized by its own particular distinction between logical form and content. For our purposes, two kinds of logical form are of primary importance. I will call these “statement form” (see II.1.2) and “predicate form” (see III.1.2). Statement logic is characterized by the notion of state-

ment form, and predicate logic, by the notion of predicate form. We should thus avoid talking about *the* logical form of a sentence as if there were a single, unambiguous form, for even when the context is rich enough to make it clear which notion is meant, this manner of speaking still leaves us with the mistaken sense that there is some unequivocal notion to be had. Even within a particular area of logic in which the notion of logical form is well defined, there is often more than one logical form of a given statement (see II.1.2.b and III.1.2.b).

3. If we plan to work with logical forms, and thus seek to remove everything from a sentence that does not belong to one of its logical forms, we need a way of representing the forms themselves. Aristotle (384–322 B.C.), the founder of formal logic, developed a *symbol system* for representing a particular class of logical forms. Representations of logical forms by means of such symbol systems are called “formulas.” Incidentally, I have already tacitly introduced Aristotle’s symbol system in Examples 1.4 and 1.5, in which particular expressions were replaced by letters (for example, ‘logician’ by *A*). Our cultural heritage is such that this procedure is almost completely transparent. We recall our algebra classes, in which all of us, with varying degrees of pleasure, learned how to work with letters in place of numbers. In algebra, as in logic, formulas represent the forms of statements: specifically, mathematical statements (see II.1.2.b, No. 9). But the relative ease with which we handle formulas should not lead us to underestimate the considerable intellectual effort involved in introducing the form/content distinction for logical forms and in developing appropriate symbol systems with which to represent them.

4. By now it should be clear what we mean by formal logic. Logic is primarily concerned with the study of logically valid inference and related notions and processes. Toward this end, there is no need to consider sentences in their entirety—it suffices to capture their logical forms. Incidentally, the name “formal logic” was not coined by Aristotle but by Immanuel Kant (1724–1804) in 1781 in his *Critique of Pure Reason* (1st German ed., p. 131; 2d ed., p. 170). Before Kant, formal logic was known

simply as “logic” or occasionally as “dialectic.” Kant believed that traditional logic could be supplemented by a new kind of logic concerned with the very content from which formal logic sought to remove itself. For reasons that need not concern us here, Kant called this new field “transcendental logic,” reserving the name “formal logic” for traditional logic as it had been known up to then.

3. Validity and Soundness

As should be clear from the preceding section, this text is concerned with questions of logical form, not content. And no formal notion is more important than that of logical validity. But our motives in seeking to understand this formal notion should be acknowledged: the reason we are so keen on validity is not just that we are interested in knowing when the truth of our premises guarantees the truth of a conclusion. Frequently, what we most want to know about a given argument or inference is, *is the conclusion true?* Knowing when an inference is valid can help us to answer this question, but it is not the whole story. Compare Example 1.1 with the following:

Example 1.7. All logicians are reptiles.
All reptiles need sleep.
*
Therefore, All logicians need sleep.

Example 1.8. All logicians are reptiles.
All reptiles are Martians.
*
All logicians are Martians.

Now, as established by our analysis of Example 1.4, it is plain that 1.1, 1.7, and 1.8 all have identical logical forms. Validity being a feature of inference form, it follows that all three are equally valid. However, it is quite

obvious that they do not succeed equally well in persuading us of the truth of their respective conclusions.

The primary difference between 1.1, 1.7, and 1.8 lies in the truth of their premises. Example 1.1 contains two true premises, whereas only one premise of 1.7 is true (the second), and both premises of 1.8 are false. Additionally, 1.7 and 1.8 differ in their conclusions. The conclusion of 1.7 is true, despite the falsehood of the first premise, and the conclusion of 1.8 is false. But to reiterate, all three inferences have exactly the same logical form; they are all equally valid. Indeed, a valid argument can have true premises and a true conclusion (1.1), false premises and a true conclusion (1.7), or false premises and a false conclusion (1.8). For that matter, an *invalid* inference can have true premises and a true conclusion! The only combination ruled out by the very notion of validity is that of a valid inference with true premises and false conclusion, since a valid inference is one for which the logical forms of premises and conclusion guarantee that whenever the former are true, so is the latter.

The usual way of drawing the distinction between 1.1 on the one hand and 1.7 and 1.8 on the other is to say that 1.1 (but neither 1.7 nor 1.8) is a *sound* inference. There are two parts to being a sound inference. A sound inference must be valid. In addition, it must have all true premises. From the understanding of validity gained in I.1, it follows that any inference that meets both conditions for soundness *must* have a true conclusion. For the fact that the inference is valid makes it truth-transferring, guaranteeing that if its premises are true, then its conclusion must also be true. But a sound inference is a valid inference with all true premises, so the truth of its conclusion may be taken as given.

Unlike validity, soundness is not a purely formal notion. Whether an argument is sound has to do not only with the logical forms of premises and conclusions, but also with whether or not the premises are true. On the other hand, soundness is not strictly a matter of logical content, either. It is worth noting that the soundness of an inference can be assessed without knowing the meanings of either the premises or the conclusion. The

only information we need in order to tell whether a given valid inference is also sound is whether the premises are true. To anticipate a notion introduced in Definition 1.2 below, all we need to know is the truth values of the premises. Thus, while soundness is not purely a matter of form, it is hardly a matter of content. Soundness is a mixed notion.

4. Statements, Primitive Statements, and Compound Statements

Truth transfer, as we have seen, is the central feature of logically valid inference (See I.1, Feature 1): When the premises of a valid inference are true, its conclusion is also true. But such talk of truth transfer would make no sense unless it was at least *possible* for the sentences in the premises and the conclusion to *be* true. It thus behooves us to distinguish those sentences that *can* be true (or false) from those that lack this feature. This is the first of three distinctions, which will ultimately lead us into statement and predicate logic.

4.1. Sentences and Statements

Are there sentences that are neither true nor false? Note that I am not here asking whether we in fact *know* a given sentence is true or false, but whether it always even makes sense to ascribe truth or falsehood to a sentence. But now observe that the first sentence of this very paragraph is an example of a sentence to which neither truth nor falsehood can be meaningfully attributed, for the sentence “Are there sentences that are neither true nor false?” is a question. In addition to questions, some sentences express norms (No turn on red!), desires (I would really like better weather tomorrow!), demands (Drop by this evening!), commands (Stand still!), and exclamations (Oh, if only Christmas could come sooner!). None of these is either true or false. This is not to say that there are not other criteria by which we evaluate such sentences—for example, as appropriate or shameless, modest or immodest, well posed or poorly posed, etc.—but they simply cannot be judged true or false. To a first approximation, it ap-

pears that only those sentences used to assert something can be true or false, and this is not what questions, commands, exclamations, and the like are for. (This is only a first approximation, since it leaves out two classes of statements, which, despite the fact that they are true or false, cannot be used to assert anything. See II.2.2.a and II.2.3.) In the part of logic concerned with truth-transferring inferences, which includes both statement and predicate logic, we must thus begin by excluding all sentences that *cannot* be either true or false. This does not mean that there is no appropriate logical treatment for such sentences, as in the logic of interrogatives or the logic of imperatives, only that we will not be dealing with them here. The following definition pins down the class of sentences with which statement and predicate logic are concerned.

Definition 1.1. A statement (or veridical sentence) is a sentence that is either true or false.

A few remarks are in order.

1. First of all, it will be helpful to introduce a standard logical convention. For this purpose, we will move right to a second definition:

Definition 1.2. In logic, “True” and “False” are called “truth values.”

As abbreviations for these values we will use T for “True” and F for “False.” One sometimes encounters other abbreviations, such as 1 for “True” and 0 for “False.” This alternative device is particularly useful in the analysis of logic gates in circuitry and in computer science in general. The expression “truth value” was introduced by the German logician Gottlob Frege (1848–1925), one of the founders of modern formal logic. Frege saw that it was possible to construe “True” and “False” as the *values* of particular functions (see chap. 4).

2. Definition 1.1 is concerned with sentences that are either true or false, and not with sentences for which it is necessarily possible for us to *determine* whether they are true or false. This distinction is extremely important in logic, as in other fields. For it is one thing to say that a sentence

has the quality of being either true or false and quite another to actually determine which it is. *Having* a truth value is a necessary (but not sufficient) condition for the possibility of our *discovering* truth or falsehood, as we cannot determine a truth value when no such truth value is present. But there are sentences for which, although we can be certain that they have a truth value, there is nonetheless no possible effective procedure for determining the value in question. Consider, for example, the statements with which we articulate natural laws. Such universal statements are purported to hold for all possible instances of a given kind, and yet there is no way to test them for all such instances. Yet despite the fundamental obstacles facing any attempt to establish their truth value, statements of natural laws are supposed to be unequivocally true or false.

3. Definition 1.1 should be understood as merely a *stipulative convention governing our use of the word ‘statement,’* and *not* as a substantive claim about all sentences that attempt to assert something. In point of fact, in ordinary usage we often assess the truth values of statements as indefinite, an assessment ruled out by Definition 1.1. According to this definition, every statement is either true or false—there is no third possibility. This insistence on the exclusivity and exhaustiveness of our two truth values is known as the “principle of the excluded middle”—in Latin, *tertium non datur* (no third is given). But for a great many sentences, which we would not hesitate to classify as statements in the ordinary sense, neither of our two acceptable truth values seems appropriate. In general, there seem to be two sorts of reasons for this failure. First, some statements are so vague as to make an ascription of truth value impossible or at least highly problematic. Consider the following:

Example 1.9. Fred is tall.

Now if Fred is over 6’6”, the statement is clearly true, but what if he is only 6’3”, 6’2”, or 5’10”? Here the problem is less a matter of *determining* the truth value of ‘Fred is tall’; it is rather a question of whether the statement even *has* truth value (see Remark 2 above). For sentences such

as 1.9, a separate brand of logic has been developed that allows for three truth values: truth, falsity, and indeterminacy. In many respects, it is an extension of the sort of logic with which we are here concerned.

A second reason for questioning the universal applicability of the principle of the excluded middle to all statements may be found in the analysis of statements about future events, including the following:

Example 1.10. The numbers in the next lottery drawing will be 2, 9, 11, 18, 21, and 35.

To claim, as the principle of the excluded middle requires, that it is *now* the case that this statement is either true or false is to commit oneself to the claim that there is already a fact of the matter as to whether this particular sequence of numbers will be drawn or not. But this is to assert that *all future events are already determined*. This view is called “determinism.” It now appears that unconditional acceptance of the principle of the excluded middle would commit us to determinism! But regardless of whether we wish to be determinists or not, it seems high-handed to force implicit acceptance of determinism on grounds of a definition in formal logic. It is for this reason that I began this remark by recommending that Definition 1.1 be viewed as merely a stipulative convention on our use of the word ‘statement’ within those branches of logic with which we are here concerned: classical statement and predicate logic. ‘Statement’ is thus defined as a technical term specific to statement and predicate logic, and as such we should expect its use to deviate from ordinary speech, just as the technical use of ‘force’ in physics deviates from everyday uses of the same term.

4. From this perspective, Definition 1.1 appears innocuous, as it merely fixes the meaning of the word ‘statement’ for purposes of statement and predicate logic. Implicit within this definition, however, we find one of the fundamental problems of logic. We are told that statements are the sorts of things that can be true or false—but what *are* these things, exactly? At present, there appear to be at least five sorts of answer, each with its own variants:

- a. Is the *assertion* made by a given statement a purely physical process or object, such as a datable acoustic blast (a speech act) or a particular, concrete inscription? If so, then a statement would be a unique, unrepeatable event or object in space and time. Such items are called “tokens.”
- b. Or perhaps a statement is an *assertion template*, an abstract pattern of sounds or written symbols, capable of numerous distinct instantiations. On this account, a statement is not a token; it is a *type*.
- c. Then again, we might choose to focus on the *act of judgment*, on the mental activity taking place within the subject who entertains a given statement. In even more distinguished vocabulary, this would mean treating the statement as a contingent attribute of consciousness. Once again, the statement becomes a spatiotemporally unique, unrepeatable event, a token.
- d. A still further possibility would be to consider the *judgment template*, the abstract type of which particular acts of judgment are the tokens.
- e. Finally, we might target the *meaning* (or *sense*) expressed by the statement, sometimes also called *the stated* or *the intended facts* (expressed by the statement in question), *the contents of judgment*, *the thought*, or *the proposition*. In which case, we must ask, to what extent and in what sense do these meanings, facts, propositions, etc. exist in relation to the four sorts of entities listed in a–d?

Needless to say, this list merely scratches the surface of the underlying problem. For even if there were deciding arguments in favor of one of the five possibilities, the attempt to elaborate the favored alternative would doubtless give rise to its own host of difficulties.

The problem of trying to determine what a statement actually is, for purposes of logic, is not a problem that arises merely out of some infelicity in the formulation of Definition 1.1; thus it is not one we can dispel by simply reformulating or perhaps omitting this definition—though in doing so the problem might be disguised. The problem here has to do with the very nature of statements, what we might call the “bearers” of truth or falsehood, and as such it is intrinsic to our subject. Since we are con-

cerned with that part of logic that deals with statements and their relations to one another, this problem is an inevitable part of our attempts to understand what logic is really about. And as the problem has turned out to be particularly difficult and multifaceted, we should acknowledge right from the start that the foundations of logic rest on shaky ground. But as is frequently (though not always!) the case in both philosophy and the sciences, this foundational instability does not preclude our pragmatic engagement with our subject. In logic, as it happens, we are able to develop very precise notions of logical form, along with techniques for checking the validity of inferences, without this fundamental problem of the nature of statements getting in the way. Indeed, the mathematical approach to logic (see IV.2) provides us with a perspective from which it is possible to abstract from the problem entirely. As far as the formal aspect of logic is concerned, we can thus leave open the question of what logic is really about. But the question remains and must eventually be addressed by anyone whose interest in logic goes beyond the purely technical. I will return to it briefly in II.3.2.

5. In what follows, we will encounter many examples of statements, all of which must be assumed to clearly satisfy Definition 1.1, despite the difficulties cited above. Merely *attempting* to do justice to these problems would leave even our simplest examples intractably long and complicated. Instead, in invoking such examples as ‘The apartment is nice,’ we will simply take them as statements in the sense of Definition 1.1. Of course, strictly speaking this sentence is extraordinarily vague. Not only does it not specify *which* apartment is meant, it also fails to restrict itself to a particular timeframe. Nor does it provide any criteria for the niceness of apartments. There will thus be numerous contexts in which the truth value of this sentence remains indeterminate. But from now on we will read this and similar examples either as having been appropriately expanded or as having been used only in appropriately restricted contexts, such that they have determinate truth value.

**Exercise 1.5.* Given Definition 1.1, which of the following sentences are statements?

1. Oh my gosh!
2. I haven't the slightest intention of considering your offer.
3. And if they haven't died yet, they're still alive today.
4. Let n be any natural number.
5. Proceed to Park Place.
6. \$200 fine for unauthorized parking.
7. You have no idea how silly you look.
8. If this goes on, it'll drive me crazy.
9. This statement is false.
10. Which of the above sentences are statements?

4.2. *Primitive and Compound Statements*

Next we divide the entire domain of statements into those that are *composed of further statements* and those that are *not composed of further statements*. The reason for introducing this distinction will become clear later on, when it emerges as an important tool in the analysis of certain classes of inferences. Consider the following example:

Example 1.11. The apartment is nice and it is fairly expensive.

This statement is composed of two separable component statements: 'The apartment is nice', and 'The apartment is fairly expensive'. They are joined into a single statement by the conjunction 'and.' Example 1.12 is similarly structured:

Example 1.12. Anthony is reading a logic book, because he finds logic incredibly exciting.

Here the two component statements are 'Anthony is reading a logic book' and 'Anthony finds logic incredibly exciting', now joined by the conjunction 'because'. We call both examples *compound* statements and say that they are composed of *primitive* statements.

Three further remarks are in order. First, it is worth noting that, in addition to such conjunctions as ‘and’ and ‘because’, the English language contains other devices that smooth the process of constructing compound statements. For instance, in 1.12, the subject of the second component statement (‘Anthony’) has been replaced by a pronoun (‘he’). Second, the conjunctions ‘and’ and ‘because’, along with other expressions by means of which primitive statements are joined to form compound statements, are called “sentential connectives” or simply “connectives” (see I.4.3 and II.1.1 below). Finally, what we call “primitive” statements are primitive only in the very specific sense that they are not composed of other *statements*. But in another sense, such statements are also compounds; like (nearly) all statements, they are composed of several words. In logic, such primitive statements are also called “atomic statements.” This label should not, however, mislead anyone into thinking that atomic statements are always really, really tiny, like the atoms of which matter is composed. The Greek root *atomos* simply means “indivisible” (although physical atoms have turned out to be divisible, a discovery which—thank God!—need not concern us here).

**Exercise 1.6.* Which of the following compound sentences are statements?

1. If you come back, may God have mercy on your soul.
2. Hopefully you’ll be doing well next time.
3. If Fred goes out with Anne, then I think Sarah will be angry.
4. It is indeed true that one candidate won the popular vote while the other won the office.
5. Someone recently asked me whether I lived in New York or Philadelphia.
6. Wait, that’s not right—or maybe it is?

4.3. *Extensional vs. Intensional Connectives*

Having distinguished primitive (or atomic) from compound statements, we now proceed to divide the connectives by which compound

statements are joined into two classes. Compound statements joined by members of the first class all have in common the fact that their truth values *depend only on the truth values of their constituent primitive statements*. We say that the truth values of such compound statements depend only on the *truth values* of their constituent primitive statements, so as to make it clear that the truth values of the compounds *do not* depend on the *meanings* of the constituent primitive statements. Consider Example 1.11: “The apartment is nice and it is fairly expensive.” To determine the truth value of the statement, it suffices to know the truth values of the two constituent primitive statements. If both are true, the compound statement is true; if one or both are false, the compound statement is false, for it simultaneously asserts the truth of both primitive statements. The connective ‘and’ is thus a member of this first class, in that it appears possible to determine the truth values of compound statements joined by this expression by reference only to the truth values of the constituent primitive statements. Neither the meanings of these statements nor any other circumstances need be taken into account.

The same does not hold for the second class of connectives. The truth values of compound statements joined by such connectives do not depend *only* on the truth values of the constituent primitive statements. Consider Example 1.12: “Anthony is reading a logic book, because he finds logic incredibly exciting.” Let us suppose that Anthony actually is reading a logic book, and let us suppose further that Anthony really does find logic incredibly exciting. Let it be given, in other words, that both of the primitive statements in Example 1.12 are true. Does this information, by itself, tell us the truth value of the compound statement? Of course, the compound statement *might* be true, if the statement accurately accounts for the motives behind Anthony’s choice of reading material. But it might as easily be false—as it would be, for example, if Anthony is doing his reading on Christmas Eve, not because he finds it so exciting (though he does), but in order to kill time while waiting to open his presents the following morning. So with connectives like ‘because’ it is usually not possi-

ble to read the truth value of the compound statement off the truth values of its constituent primitive statements. More information is needed.

The distinction between these two classes of connectives is expressed in the following definition:

Definition 1.3 1. *Extensional* (or *truth-functional*) connectives are such that the truth values of compound statements joined by these connectives are completely determined by the truth values of the constituent primitive statements.

2. *Intensional* connectives are such that the truth values of compound statements joined by these connectives are not determined by the truth values of the constituent primitive statements alone.

***Exercise 1.7.** Which of the following connectives are truth-functional: ‘or’, ‘due to the fact that’, and ‘if . . . then’.

Definition 1.3 introduces some novel terminology. The extensional connectives are called “truth-functional” because the truth values of the compound statements joined by them can be viewed as a *function* of the truth values of the constituent primitive statements. In general, to say that y is a function of x means that the value of variable y is unequivocally determined by the value of variable x . “Intensional” comes from “intension,” which refers to the meaning of a sentence, and “extensional” means simply “independent of the meaning of a sentence.”

Once again, proper appreciation of the above distinctions will have to await our development of statement logic in the following chapter.

5. Review

The following will suffice as a provisional summary of this introduction to the central issue of formal logic. Formal logic is primarily concerned with valid inferences—in other words, with inferences in which the logical forms of premises and conclusion force us to accept the latter

once we have accepted the former. More generally, formal logic deals with those features of statements, and of the relationships between statements, that may be said to hold *by virtue of the logical forms* of the statements in question. Valid inference is an example of such a relationship, and logical truth, about which we will learn more later, will be an example of such a feature. Our discussion has presupposed the sort of prior understanding required in order to evaluate basic inferences as valid or invalid, several examples about which were analyzed. Over the course of several sequential steps, I motivated the introduction of a notion of logical form. Finally, we paved the way for a proper treatment of statement and predicate logic with the following distinctions: statements versus sentences in general; compound versus primitive statements; and extensional versus intensional connectives. We turn now to a detailed discussion of the extensional connectives and their properties.