CONTENTS

CHAPTER 1

Is Logic Boring and Pointless?

1.1 Logic in Practice, Logic in Theory

Logic is easy in specific cases, but difficult in general.

If I tell you that all cats are mammals and that all mammals are animals, then you conclude that all cats are animals. If another time I tell you that my cat is always asleep at 4 o'clock in the afternoon, and you then notice that it *is* 4 o'clock, then you conclude that my cat is asleep. If instead you see my cat walking around and plainly not asleep, then you conclude that it is not 4 o'clock.

That is logic.

This is simple. These conclusions are *obvious*. Suppose, though, that someone doubts you. He asks, "How do you *know* it follows that all cats are animals? How can you be *certain* that the cat is sleeping or that the time is 4 o'clock?" You would hardly know what to say. The fundamental principles of logic seem so straightforward and intuitive that it is unclear how to explain them in terms of something more readily comprehended. You would assume that the skeptic in some way misunderstands the language. You might even repeat the premises to him, slower and louder.

The ancient philosopher Sextus Empiricus, writing in the second century CE and commenting on the work of the Stoic logician Chrysippus (279–206 BCE), suggested that even dogs understand these principles:

> [Chrysippus] declares that the dog makes use of the fifth complex indemonstrable syllogism when, on arriving at a spot where three ways meet. . ., after smelling at the two roads by which the quarry did not pass, he rushes off at once by the third without stopping to smell. For, says the old writer, the dog implicitly reasons thus: "The animal went either by this road, or by

> that, or by the other: but it did not go by this or that, therefore he went the other way." (Floridi 1997, 35)

There really does seem to be something instinctive about the principles of logic. Every time you search for your lost keys by retracing your steps, you are applying those principles. You say to yourself, "I know I had the keys when I left the house. I then visited locations A, B, and C, and I could only have left the keys at a location that I visited; therefore, I left the keys at one of locations A, B, or C." Of course, you never really pause to spell out the steps of the argument, and that is precisely the point. You process the basic logic of the situation so automatically you are hardly aware that you have reasoned at all.

The extreme naturalness of logical reasoning was noted by John Venn, in his 1881 book *Symbolic Logic*:

> It may almost be doubted whether any human being, provided he had received a good general education, was ever seriously baffled in any problem, either of conduct or of thought . . . by what could strictly be called merely a logical difficulty. It is not implied in saying this, that there are not myriads of fallacies abroad which the rules of Logic can detect and disperse. . . . The question is rather this: Do we ever fail to get at a conclusion, when we have the data perfectly clearly before us, not from prejudice or oversight but from sheer inability to see our way through a train of merely logical reasoning? . . . The collection of our data may be tedious, but the steps of inference from them are mostly very simple. (Venn 1881, xix)

As Venn notes, people do commit logical errors. In my role as a mathematics educator, I encounter such errors all the time. For example, it is common for students to treat a statement of the form, "If p, then q," as though it were logically equivalent to the statement, "If q, then p." Rarely, though, is the student confused about abstract principles. If you point out to him that the statement "If Spot is a normal dog, then Spot has four legs" is true, while the statement "If Spot has four legs, then Spot is a normal dog" is false, he immediately understands your point. He does not argue that *he* is the one who is thinking clearly, while *you* are the one who is confused. Such errors occur simply because mathematicians deal with complex statements about abstract objects, and, in the heat of the moment, students often find them difficult to parse.

Returning to my opening examples, the conclusions not only followed from the premises, they followed in a particular way. Consider a contrasting example. Recently I was preparing dinner for a large group of friends. I was in the kitchen, with the exhaust fan, the sink, and the television all going. It was *noisy*. Suddenly, one of my cats came barreling in, her paws struggling for purchase on the smooth kitchen floor. She darted down the basement stairs. I reasoned, “My cat only panics like that when she hears strangers on the patio. I’ll bet my guests have arrived, but I did not notice the doorbell amidst all the racket.” I went to the door and found that I was right.

I came to this conclusion because I knew many instances of my cat panicking at the sound of strangers on the patio, and no instances of her panicking in that manner in response to any other stimulus. I reasoned from empirical facts and extensive personal experience to the conclusion that my guests had arrived. Philosophers would say that my reasoning was *inductive*, from Latin words that translate roughly as “to lead into.” In this case, I was led to a general conclusion from my experiences in a few specific cases. This style of reasoning is common in science, where our confidence in a theory’s correctness grows each time it accurately predicts the outcome of an experiment.

Our opening examples were not of that sort. From “All cats are mammals” and “All mammals are animals,” we concluded that “All cats are animals,” but this conclusion in no way followed from anything we know about cats, mammals, or animals. From “All cats are green” and “All green things are plants,” it follows that “All cats are plants,” though in this case, all three statements are false. Our notions of what follows logically from what are unrelated to the facts of the world. Instead, they are related in some manner to the way we use language and to the grammatical structure of the assertions involved. We could say “All As are Bs, all Bs are Cs, and therefore, all As are Cs” without causing controversy.

This sort of reasoning is known as *deductive*, from Latin words meaning “to lead down from.” It is the primary sort of reasoning employed in mathematics. Deductive reasoning seems to have a certainty about it that inductive reasoning lacks. My conclusion that my guests had arrived given the evidence of my panicked cat was perfectly reasonable. However, it might have been that my cat had been scared by something else, or was just being weird for some reason. My conclusion *might* have been wrong. But if the statements “All cats are mammals” and “All mammals are animals” are both true, then it simply *must* be true that “All cats are animals.” Period. End of discussion.

This all seems sufficiently straightforward. The difficulty comes when trying to formalize our intuitive notions. Can we write down a general set of rules to tell us what follows from what?

We have seen that logical inferences are closely related to language, and, indeed, “logic” comes to us from the Greek “logos,” meaning “word.” In natural languages—English, French, German, and so forth—there are many types of words. There are nouns, which we can take to represent objects, and verbs, which generally describe what the nouns are doing to each other. There are adjectives to supply additional information about the nouns, prepositions to describe relationships among them, and adverbs to tell us more about the verbs.

Then there are other words whose function is to establish logical relationships among the component clauses of a complex sentence. Philosophers refer to these words as *logical constants*. In English, we use such words and constructions as "not," "and," "or," and "if–then" as logical constants. You come to understand the meanings of these words by understanding the effect they have on the clauses to which they are connected.

For example, if I say, "On Tuesday, I ate cookies, and I ate cake," the role of "and" is to tell you that I ate both cookies and cake on Tuesday. If you later discover I only ate one of them, or neither of them, you would think I had said something false. In this context, we come to understand what "and" means by understanding the truth conditions it imposes on the sentence whose clauses it is joining. Moreover, "and" plays this role regardless of the content of the clauses on either side of it. That is why it is called a "logical constant."

This is progress toward our goal of having general rules for telling us what follows from what. If we let p and q represent simple assertions, then we can say that from the sentence "p and q" we can fairly conclude that p and q are both true individually. Given some familiarity with standard English usage, we can quickly write down other such rules:

- The statement "not p" has the opposite truth value from p.
- Given "p or q," and "not p," we can conclude that q is true.
- Given "If p, then q," and "p," we can conclude that q is true.

There is much that could be added to this list, of course. For the moment, however, the main point is that this logic business does not seem very complicated at all. Writing down logical rules involves nothing more than understanding what words mean, and you hardly need a degree in philosophy for *that*.

Matters are not always so simple, however. If I am at a restaurant, the server might ask me whether I want french fries or mashed potatoes with my dinner. Later he might ask me whether I want coffee or dessert. In the first instance, it is understood that I am to choose only one of french fries or mashed potatoes, while in the second it would be acceptable to have both coffee *and* dessert. What, then, should the rule be for statements of the form "p or q"? If p and q are both true individually, should "p or q" be deemed to be true? Or is it false? It would seem there is no rule that covers all contexts.

And how are we to handle conditional statements, by which I mean statements of the form "If p, then q"? If p is true by itself, and q is false by itself, then "If p, then q" should be considered false. *That* much is clear.

But what if p and q are both true? Should we automatically declare "If p, then q" to be true in this case? That seems reasonable for mathematical statements: "If x and y are even numbers, then $x+y$ is even as well," for example. In contrast, what am I to make of the statement, "If I am not a cat, then I am not a dog"? Both parts are true by themselves, but the sentence as a whole does not seem to be true. In everyday usage, we normally take it for granted that the two parts of a conditional statement are relevant to each other, but it is unclear how to capture a relevance requirement in a logical system.

Natural languages have many other attributes that make logical analysis very difficult. They contain statements that are vague or ambiguous. Some statements are *indexical*, which is to say that their meaning depends on the context. For example, the meaning of "I am hungry" changes, depending on the speaker. The truth or falsity of a statement often depends on more than just its grammatical structure. For example, the statements "If my cat did not eat the tuna, then someone else did" and "If my cat had not eaten the tuna, then someone else would have" have very different meanings, though we might naively interpret both as having the abstract form "If p, then q."

It would seem that trying to capture the logical rules implicit in everyday language is not so simple after all.

Seeking respite from such travails, logicians prefer instead to work with formal languages. By a "formal language," I mean a language the logician simply invents for her own purposes. The logician therefore has complete control over what counts as a proper assertion, and she can devise strict rules for determining the correctness of proposed inferences. There is no vagueness and no ambiguity. For logicians, the move from a natural to a formal language produces a calming effect, similar to when the kids are out for a few hours and blessed quiet descends on the house.

In crafting her language, the logician might begin by inventing symbols to represent basic sentences. Other symbols are then devised to denote familiar connectives, like "and," "or," and "if–then;" and still more symbols are introduced to denote various sorts of entailments and implications. As a result, simple assertions can be made to look complex. For example, our inference that all cats are animals from the assumptions that all cats are mammals and all mammals are animals, might end up like this:

$$(\forall x\ Cx \to Mx) \wedge (\forall x\ Mx \to Ax) \models (\forall x\ Cx \to Ax).$$

You should interpret "Cx," "Mx," and "Ax" to mean, respectively, that x is a cat, mammal, or animal. The upside down A is an abbreviation of "for all," the arrow means "if–then," and the vertical wedge means "and."

The symbol that looks like the Greek letter pi on its side denotes entailment. Thus, translated back into English we have, "The assumptions that for all x, if x is a cat, then x is a mammal, and for all x, if x is a mammal, then x is an animal, entail the conclusion that for all x, if x is a cat, then x is an animal."

Practitioners of formal logic are fond of this sort of thing. A statement as simple as "My cat is furry" might be rendered thus:

$$(\exists x)\,(Jx \wedge (\forall y)\,(Jy \rightarrow (y = x) \wedge Fx))\,.$$

In English, this collection of symbols means: "There exists an x such that x is Jason's cat, and if y is anything else that is Jason's cat, then y is the same as x, and y is furry." Where you might see a simple statement of fact about my cat, a logician sees a complex existential assertion involving conditional statements and conjunctions. This, from a sentence containing neither the word "and" nor "if–then." It would seem that a difficult logical structure of language lurks beneath its grammatical structure.

The relationship of the formal language to natural language is like that of a laboratory experiment to the real world. Scientists contrive controlled scenarios in which a few variables can be studied in isolation from others. They then hope that they have chosen the really important variables, so their results will be applicable to reality. Likewise, the logician hopes that the formal language captures those aspects of natural language that are relevant to reasoning, even though she knows subtle aspects of the natural language will inevitably be lost in her formalization.

1.2 *Enter the Philosophers*

The translation of simple, natural-language sentences into difficult symbolic ones can be a tedious affair, but the worst is still to come. Once the philosophers learn of your project, they will want a piece of the action, and God help you when *that* happens. Philosophers have investigated, *minutely*, all of the central notions on which logic relies. Through their investigations, they have discovered the only thing philosophers ever discover: that everyday notions used without incident in normal social interactions become murky when closely analyzed.

For example, most elementary textbooks will tell you that the fundamental unit in logic, comparable to atoms in physics or prime numbers in arithmetic, is "the proposition." If we ask, "What sort of thing is it that can rightly be described as either true or false?" the answer is, "A proposition." It is gibberish to say, "This vegetable is true" or "This color is false," but it makes perfect sense to say, "This proposition is either true or false."

But what *are* propositions?

One possibility is that a proposition is just the same thing as a declarative sentence. This seems plausible. In a conversation, we might say, "What you just said? That's so true!" when what the person just said is actually a sentence. High school students often take examinations in which they are asked to mark each of a list of sentences as being either true or false. So, maybe the concept of "proposition" does not really add very much, and we should just talk directly about sentences instead.

The problem is that the same sentence can mean different things in different contexts. When I say, "I have a cat," I am not expressing the same proposition as you are when you utter those same words. However, two different sentences might express the same proposition. In France, I would say, "J'ai une chat" instead of "I have a cat," but the same proposition has been expressed. Some sentences do not seem to express any proposition at all. "It is raining" is a perfectly fine sentence, but until we contextualize it to a time and a place, we cannot assign it a truth value. In light of these considerations, it seems accurate to say that we use declarative sentences to express propositions, but that the sentences are not themselves propositions. There are concepts of some sort to which sentences point, and *those* are the propositions.

Moreover, propositions do not just get stacked up into written arguments so that other propositions might be drawn as conclusions from them. They have another existence as beliefs in a person's mind. When you believe something about the world, what kind of thing is it that you actually believe? A proposition, that's what! However, it does not seem right to say that the thing you believe is a sentence, as though you cannot have beliefs unless you have first summoned forth sentences that express them. My cat has beliefs about the world, but, clever though she is, I doubt that she can express those beliefs in sentences.

So the question persists. What *are* these things we call "propositions"? Are they just the meanings of sentences? Can we define "proposition" as "what a sentence means"? Perhaps, but does this really help us understand what is going on? Meaning is itself a very difficult concept, as pointed out by philosopher A. C. Grayling in a discussion of this very point:

> Suppose I am teaching a foreign friend English, a language of which he is wholly ignorant; and suppose I point to a table and utter the word 'table.' What settles it for him that I intend him to understand the object taken as a whole? Why should he not take me as pointing out to him the colour, or the texture, or the stuff of which the object is made? Imagine my pointing at the table-top and saying 'glossy.' Why should he not understand me as naming

> the object as a whole, rather than the style of its finish? At what is apparently the simplest level of demonstratively linking a name with the object it is supposed to 'mean,' then, there are puzzling difficulties. (Grayling 1982, 36)

If meaning is difficult even in this simple case, then how much more difficult is it when we speak of the meaning of a whole sentence? For example, how is understanding the meaning of a sentence different from just understanding the proposition it expresses?

It is at this point, when most people find their eyes glazing, that the philosophers start to get *really* interested. Their chief weapon in the fight against vagueness and imprecision is the drawing of subtle distinctions, and the literature in this area offers plenty of them: between sentences, statements, and propositions; between sense, meaning, and reference; between the intension and the extension of a term. Not to mention what is potentially the most important distinction of all: between realism and nominalism with respect to abstract objects. You see, if you take the view that there are these spooky, ill-defined propositions floating around just waiting to be gestured at by sentences, then you sure seem to be suggesting that abstract objects actually exist. That makes you a realist. Against you are the nominalists, who regard abstract objects as useful fictions that humans devise for their own purposes. (Does the number three actually exist as an object by itself? Or is "three" just a name we use to describe what is common among all collections of three physical objects?) This particular dispute has raged for centuries, and I assure you that the rival camps see this question as *very* important.

Do you see what happened? We asked, in perfect innocence, what propositions were, and just a few paragraphs later, we were mired in deep questions of ontology and metaphysics. For heaven's sake.

Let us put these niceties aside. Assume for the moment that we have arrived at a coherent account of "proposition." What does it mean for this proposition to be "true"?

Any nonphilosopher would say that the true propositions are the ones in accord with the facts. We have facts on one side, true propositions on the other, and for every true proposition, there is a corresponding fact that makes it true. What could be simpler?

We could retort, however, that this approach is *too* simple. A philosopher might say, "Yes, thank you, I *know* that truth is about correspondence with facts in some vague way, but that is unhelpful. I need to understand the process by which a proposition is paired with the fact to which it corresponds. If I asked, 'What caused this patient's death?' you would no doubt reply, 'He died because his heart stopped,' thinking you had thereby said something informative. But the question, *obviously*, is what caused his heart to stop. Likewise, the question for those claiming

Figure 1.1. The empirical fact corresponding to the proposition, "My cat is watching me type this."

that truth is about correspondence with fact is to explain the nature of this correspondence, and good luck with that."

How *do* we go unerringly from the true proposition to its corresponding fact? Correspondence seems straightforward when considering simple assertions. "My cat is watching me type this" is true because of a certain empirical fact, depicted in Figure 1.1. Matters are far less straightforward when discussing complex statements. What is the fact of the world corresponding to "If my cat had not broken her leg, I would have spent Saturday either reading a book or watching television, instead of rushing her to the veterinarian"? It would seem that facts can be rather complex. To make the correspondence theory work, I would first need an account of what facts are and then an account of the manner in which the pairing of true propositions with facts is achieved. Neither of these accounts is readily forthcoming.

Other types of sentences cause problems as well. What fact of the world corresponds to "There are no unicorns"? Perhaps the relevant fact is found by restating the sentence in the equivalent form: "Everything is a non-unicorn," but, among other problems, this suggests that a sentence that certainly appears to be about unicorns is actually about literally everything except unicorns. Similar problems could be adduced for disjunctions (or-statements), counterfactuals, statements about the past, and statements about abstract entities (like $2+2=4$). In each case, it is not straightforward to identify the piece of reality to which the proposition corresponds.

The more you think about it, the more difficult it becomes to pin down the correspondence relation that is said to obtain between true propositions and facts. Propositions are abstract entities, some notion of which resides in our heads. Facts are about physical objects that exist out there in the world. "Correspondence" implies some sort of isomorphism between these radically different realms. How can that be?

Perhaps you think the solution is as follows. We begin by identifying certain simple, basic facts. These correspond straightforwardly with simple propositions, by which we mean propositions with no logical structure to them. The orange cat on my sofa corresponds simply to the proposition, "My cat is orange." The facts corresponding to more complex statements are then found by breaking the statements down into the logical simples out of which they are made. Done!

The philosophers have a name for this approach, which is never a good sign. It is called "logical atomism," the idea being that these logical simples are like the atoms out of which chemical substances are made. At various times, this approach has been defended by giants like Bertrand Russell, Ludwig Wittgenstein, and Rudolf Carnap. Nowadays, however, the notion has fallen on hard times, for reasons you have probably already guessed. Those logical simples have proved surprisingly elusive, and no one has managed to supply a helpful account of them.

It would seem that the correspondence relation is so murky and complex that we might reasonably wonder whether it is actually helpful in elucidating the nature of truth. The main argument in favor of the correspondence theory (really, the only argument) is its agreement with common sense. In daily life, it sure feels like we assess truth first by understanding a sentence's meaning and then by comparing it with relevant facts. Philosophers, though, take special delight in refuting common sense. Tell a philosopher that an idea is intuitively obvious, and he will quickly retort that, so sorry, it is incoherent nevertheless.

At this point, we might think that our whole model is wrong. We have been acting as though we have the world of propositions *over here*, and then separately from that, there is an objective reality *over there*. This objective reality comes equipped with facts, and in some vague way, it is these facts that make propositions true. The relation between fact and proposition is said to be one of correspondence, but we encountered difficulty spelling out the nature of this relation.

There are other possibilities. Maybe it is not facts (whatever *they* are) that make propositions true, but rather other propositions. That is, we could say that a proposition is true when it coheres with other propositions that are already accepted. Defenders of this view argue that the relation of coherence is more readily described than that of correspondence. Or maybe the whole concept of truth is just redundant.

After all, what is the difference between saying, "Proposition *p* is true" and just asserting *p* in the first place? In this view, stating that a proposition is true is different from stating that an apple is red. The latter case attributes a property to an object, while the former does not. These are called the "coherence" and "redundancy" theories of truth, respectively. They have their defenders, as do several other theories I have chosen, because I want people to keep reading my book, to omit.

Mighty treatises and mountains of journal articles have been written on each of these matters, and believe me when I tell you, they do not make for light reading. Nothing to relax with before bed in *that* charming little ocean of verbiage. Perhaps, though, we are justified in ignoring this literature. Just as I can drive a car without knowing how it works, so, too, can I use notions like "proposition" and "truth" without a proper philosophical account.

Sadly, though, we are just getting started. Once you start asking philosophical questions about logic, it is impossible to stop. Do the laws of logic exist by necessity, are they just arbitrary consequences of the way we define words, or are they empirical facts discovered through investigation and experiment? Should logicians be seeking the one true logic that applies always and everywhere? Or are systems of logic more like systems of geometry: useful or not useful in different contexts, but not correct or incorrect in any absolute sense? Should "true" and "false" be regarded as the only truth values? Some statements are vague, after all, and therefore do not fit comfortably into a binary conception of truth. To accommodate this fact, perhaps we should countenance truth-value gaps, by which I mean propositions that are neither true nor false. Perhaps we want a third truth value, "neutral," which applies to statements that are vague. Maybe we should countenance the possibility of truth-value gluts, as when we find that a proposition is true in one sense but false in another. Maybe "both true and false" ought to be an option. How should we handle different modalities? Some propositions are possibly true, as when I say, "Tomorrow I will go to the park," while some are necessarily true, like "Two plus two equals four." Should our system of logic reflect this difference?

No point about logic is too clear and simple to avoid the complexifications of philosophers. Delve into this literature, and you will encounter careful discussions of the distinction between the philosophy of logic and philosophical logic. You will see people get the vapors over the problem of assigning a truth value to the sentence: "This very sentence is false." (If it is false, then what it asserts is actually the case, so it is true. But if it is true, then what it asserts must be the case, which makes it false!) You will meet people of unquestioned brilliance and sagacity thinking that it is an insightful commentary on the nature of truth to observe that "Snow is white" is true if and only if snow is white.

And you will encounter people who argue, in perfect seriousness, that some contradictions are true. You read that right. No idea is so daft that some philosopher has not floated it.

Folks, *this* is what awaits you if you choose to study logic. Textbooks on formal or mathematical logic appear to be written in hieroglyphics. They are also often disappointing, in that after translating the symbols into English, it is common to find that something trivial has been asserted. Texts on philosophical logic are well-nigh unreadable, and often leave you with the uncomfortable feeling you know less after reading them than you did before. Whatever the precise subject matter, if the word "logic" is in the title, it is likely not the sort of book most people would enjoy reading. Taking a course in the subject would rank low on almost anyone's bucket list. "Recreational logic," in this view, must simply be dismissed as an oxymoron.

Which is a pity, since deducing the logical consequences of a set of premises can be surprisingly enjoyable.

1.3 Notes and Further Reading

The philosophy of logic is an inherently difficult subject, and even its introductory texts make for heavy reading. I found the books by Haack (1978), Grayling (1982), and Read (1995) to be especially helpful. Sarcasm aside, I found all three books to be fascinating, even if I was occasionally unpersuaded of the importance of some of the more esoteric discussions. While it is true than one can drive a car without understanding how a car works, I am certainly happy there are people out there who *do* understand how cars work.

The correspondence theory of truth has recently been the subject of a book-length defense by Joshua Rasmussen (2014). In particular, he provides a detailed discussion of the sentence, "There are no unicorns." Both this book and the references contained therein will be helpful to anyone wanting to look into this area.

The view that some contradictions are true is referred to as *dialetheism*. It is definitely a fringe view among logicians generally, but it is taken seriously and has enthusiastic defenders. At one point, I considered trying to devise some dialetheistic logic puzzles for this book, but ultimately I found the idea just a little too unnatural to wrap my head around. At any rate, philosopher Graham Priest has been the most eloquent champion of the idea, and you can check out his books *In Contradiction* (2006) and *Doubt Truth to Be a Liar* (2008a) for more information.

INDEX