CONTENTS

INTRODUCTION

Many people believe that philosophy is abstract, esoteric, and practically irrelevant. It might be fun to think about in your spare time, but it isn't useful for anything else. This common perception applies equally to *epistemology*—the theory of knowledge. If you ask, "How do you know you're not dreaming right now?" at a cocktail party, that could be fun, but if you ask it in the middle of a job interview, it will seem a little odd, to say the least. That question just isn't *practical*. And that is how people perceive philosophy, in general. It's fun, but it isn't practical. Hopefully, this book will demonstrate that this is very far from the truth. Many of the problems that we face in our everyday lives are actually philosophical problems; they just aren't *recognized* as philosophical problems.

Consider just one example. In the modern world, most of our knowledge comes from other people and through books, magazines, broadcasts, and podcasts. Of course, we know that there is a lot of deception in the world. This forces us to ask the question: *Who should I trust?* Well, when we trust other people, what is our goal? We trust people for lots of different reasons—to cooperate, to build friendships, to gain their trust in return. But sometimes, we trust people simply because there are things we want to *know*. When that is our goal, we should trust people when, and only when, we will get knowledge from them. If that is our goal, then in order to know who to trust, we need to answer another question: *When does trusting another person's testimony produce knowledge?* That is a philosophical question. It's practical, and it's also philosophical. This book is about those kinds of questions.

Each chapter of this book will begin with a true story. As we will see, each story poses an epistemological problem—a problem about knowledge. In Chapter 1, we begin with the story of Bernie Madoff, who deceived everyone on Wall Street and defrauded thousands of people from their life savings. Madoff's story illustrates exactly what

we want to avoid when we trust others. How can we do that? Under what conditions do we *know* that a person is telling us the truth? That is the question addressed in Chapter 1. In Chapter 2, we begin with the economic crisis of the late 1970s, in which both unemployment and inflation were on the rise. President Jimmy Carter needed to find someone who could fix the economy. He needed an expert. Since Carter was not himself an expert in economics, how could he know who the experts were? That is the question addressed in Chapter 2.

Chapter 3 begins with the extraordinary friendship between Francis Collins and Christopher Hitchens. Collins, a renowned scientist, was also an evangelical Christian, while Hitchens was an avowed atheist. Each of them was highly intelligent and well-informed, yet they completely disagreed about religion. How should they each respond to their disagreement? Should it lead them to revise their own beliefs? That is the question addressed in Chapter 3. In Chapter 4, we begin with a story about cognitive bias in a murder investigation. This story illustrates a general truth about all human beings—we all suffer from many cognitive biases. In fact, the extent of our biases is so great that it casts doubt on our ability to know the truth in many cases. What should we infer from this? That is the subject of Chapter 4.

In Chapter 5, we begin with the story of political polarization, which is on the rise today. Many people form their political beliefs by trusting the people in their community or in-group. In an age of polarization, this seems to imply that many people must be forming false beliefs in this way. How do you know that it isn't *you* who are forming false beliefs in this way? That is the subject of Chapter 5. Chapter 6 starts with the story of an American propaganda campaign in several countries in eastern Europe and central Asia. The existence of propaganda threatens our ability to acquire knowledge from news outlets. Chapter 6 also investigates what is called "nudging" and "Big Data." Like propaganda, these are ways in which our beliefs can be manipulated. Do they all deprive us of knowledge, and if so, then why?

Chapter 7 begins with some recent successes of artificial intelligence. AlphaGo—the computer program that plays the strategy

game Go—stunned the world by beating the world champion Go player, and AlphaFold predicted the structures of many proteins, which even the best biologists had been unable to do. Can AI give us scientific knowledge? Can it do this on its own, independently of human researchers? That is the question addressed in Chapter 7. Finally, in Chapter 8, we begin with a conspiracy theory. In the wake of Hurricane Katrina, many people believed that the levees around New Orleans had been bombed deliberately to destroy black neighborhoods. Is it always irrational to believe a conspiracy theory like this? Why or why not? That is the subject of Chapter 8.

Throughout this book, you will see that our everyday lives pose epistemological problems. You can use the concepts and theories developed in epistemology to solve these problems for yourself. They will also give you the tools you need to solve new problems in life as they arise. As you will soon see, epistemology is a very useful, practical discipline.

In the remainder of this Introduction, I will define some key terms, as I will use them throughout this book. I will also explain one key assumption that I will make. I will not assume any particular theory of knowledge. As we will see, different theories sometimes give different answers to the questions we discuss, so we will not assume any particular theory at the outset. However, in order to avoid confusion, I will clarify how I use some key terms, as well as one basic assumption that I will make.

First, knowledge requires that you have a *true belief*. You can know that the earth is round because it is true that the earth is round. By contrast, no one knows that the earth is flat, simply because it is false that the earth is flat. You cannot know what is false. That is part of our ordinary concept of knowledge—it is part of how we use the term "knows," and I will follow that ordinary usage in this book.

Second, while true belief is necessary for knowledge, it is not sufficient for knowledge. In other words, it is possible to have a true belief without *knowing* that your belief is true. Suppose that you think of a number between one and ten, and you ask me to guess the number. I guess that you are thinking of the number seven. It just so happens that you *are* thinking of the number seven. Then I have a true

belief—I believe that you are thinking of the number seven, and it is *true* that you are thinking of the number seven, so my belief is true. But do I *know* that you are thinking of the number seven? No, I don't know. I just made a lucky guess. So, just having a true belief is not yet knowledge. Knowledge requires something more.

What more is required to have knowledge? As the example above illustrates, if you were just lucky that your belief is true, then you don't really know that it is true. Thus, we can infer that knowledge is true belief that is *not just true by accident.* What does that mean? Here is one way to understand this.[1] Some ways of forming beliefs are *reliable*—they usually lead to true beliefs. For example, if you have typical adult human vision, then using your vision to form beliefs about what is right in front of you is a very reliable way of forming beliefs. That process will lead to true beliefs most of the time because your vision is very reliable, at least at very close range. Thus, typical adult vision, at close range, is a reliable way of forming beliefs.

Remember that knowledge is true belief that is not just true by accident. Suppose I have typical adult vision, and I use my vision to form a belief about what is right in front of me. In those circumstances, if I form a true belief about what is in front of me, *it will not be an accident that I form a true belief.* That is because I form my belief by using a process that is very reliable. It is a process that leads to true beliefs most of the time. Thus, it will not be an accident if that process leads me to form a true belief since that is what it tends to do. When you use your eyes to determine what is right in front of you, you will usually form a true belief, and so it will not be an accident if you form a true belief in that way. Now suppose that I use my vision to form a belief about what is right in front of me, and I thereby form a true belief. Now I have a true belief *that is not just true by accident*; thus, I have *knowledge*. Finally, here is what we can infer from this: *if you form*

1. In what follows, I will use the terminology of a particular theory of knowledge, which is called *reliabilism.* However, this is only to illustrate the fact that knowledge can be *fallible.* As I will explain, it is that latter claim—that knowledge can be fallible—that I will assume in this book. I will not be assuming the reliabilist theory of knowledge.

a belief in a reliable way—a way that usually leads to true beliefs—and you thereby form a true belief, then you have knowledge.

I have said this in order to explain an important assumption I will make in this book. I will assume that *knowledge is fallible.* To understand this, consider again a case of using my vision to form a belief about what is right in front of me. Suppose that I look directly in front of me, I see a laptop computer in front of me, and I form the belief that there is a laptop computer in front of me. In this situation, I have used a reliable process—my vision—to form a true belief. Thus, I *know* that there is a laptop computer in front of me. However, the process that I used to form this belief—trusting my vision—is not an *infallible* process. Even that process can lead to error. On this particular occasion, it did *not* lead to error, but it is always *possible* that it led to error. Consider just a couple of ways in which this could happen. If I had been drugged with a hallucinogenic drug, then I might hallucinate a laptop computer, even if there is not one there. Alternatively, someone in the physics department might decide, as a prank, to project a hologram of a laptop computer onto my desk, and I cannot tell the difference just by looking at it. These are unusual examples, but they illustrate the point: even adult human vision at close range is *fallible*—it can lead to false beliefs.

However, despite its fallibility, my vision often gives me knowledge. In ordinary circumstances, when my vision is working properly and it produces a true belief, I *know* that my belief is true. Since I formed my belief through a process that usually leads to true beliefs, it is not an accident that my belief is true. Thus, I have a true belief that is not just true by accident, and that is knowledge. In summary then, if a way of forming beliefs is *reliable*—it usually leads to true beliefs—then it can give us knowledge, even if it is *fallible.* A way of forming beliefs need not be infallible in order to give us knowledge. It only needs to be reliable. This view is called *fallibilism* about knowledge, and in this book, I will assume that it is true. It is certainly true with respect to *testimony.* Testimony is fallible, but it often gives us knowledge. The question is *how* and *when* testimony gives us knowledge. It is to that question that we now turn.