THINKING IT THROUGH

AN INTRODUCTION TO CONTEMPORARY PHILOSOPHY

Kwame Anthony Appiah

Thinking It Through

AN INTRODUCTION TO CONTEMPORARY PHILOSOPHY

Kwame Anthony Appiah

OXFORD
UNIVERSITY PRESS
2003

CONTENTS

PREFACE

You learn a lot about your subject when you set out to introduce the range of it to people who are approaching it for the first time. That is a good part of the reason I set out to write an introduction to contemporary philosophy. After a while, as you do the detailed work of professional research, you risk losing sight of the forest for the trees. Stepping back for a bit, to think again about the shape of the subject and where your own work fits into it, allows you not just to rediscover connections but also to make new ones. That is why undergraduate teaching is so invigorating.

What I have tried to write is a reliable and systematic introduction to the central questions of current philosophical interest in the English-speaking world. (I have also pursued some less mainstream questions because I think they should be more mainstream!) A philosophy textbook can't be a record of current answers to the central questions, because philosophy is as much about deepening our understanding of a question as it is about finding an answer. So my task has been to prepare the reader to enter into contemporary debates by delineating the conceptual territory within which the many answers currently in play are located. I hope I have succeeded in making it possible for a newcomer to navigate that territory and that I have also made the navigation seem engaging, for that will mean that some of my readers will want to read more deeply in the subject. An introduction can be the beginning of a lifelong romance.

I find I have now taught philosophy on three continents, and it is astonishing how the same questions arise in such culturally disparate circumstances. I am grateful to all of my students, in Ghana, in England, and in the United States: Almost every one of them has taught me a new argument or—what is much the same—shown me an old one in a new light. This book is dedicated to them.

INTRODUCTION

A Few Preliminaries

People come to philosophy by many different routes. The physicist Schrödinger, who developed some of the key concepts of modern quantum theory, was drawn into philosophy by the profoundly puzzling nature of the world he and others discovered when they started to examine things on the scale of the atom. One of my friends came to philosophy when, as a teenager, he was first developing adult relationships of friendship and love. He was perplexed about how easy it was to think you understood somebody and then discover that you had not understood her at all. This led him to wonder whether we ever really know what is going on in other people's minds. And many people come to philosophy when they are trying, as we say, to "find themselves": to make sense of their lives and to decide who they are.

If, for these or any other reasons, you come to have an interest in philosophy, it is natural to turn to the works of great philosophers. But for most people the content of these works is rather a shock. Instead of offering direct answers to these questions—What is physical reality really like? Can we ever be sure we know what other people are thinking? Who am I?—a philosopher is likely to start with questions that seem to him or her more basic than these . . . but which may seem to others far less interesting. Instead of beginning by asking what we can know about other people's thoughts, a philosopher is likely to start by asking what it is to know anything at all—thus beginning with epistemology, which is the philosophical examination of the nature of knowledge. Despite the natural disappointment it produces,

I think that starting with these fundamental questions makes sense. Let me suggest an image that might help you to see why.

Imagine you are lost in a large old city in Africa or Asia or Europe. Every way you turn there is interest and excitement. But you'd like to know where you are. The trouble is that just when you think you have found your way out of one maze of alleys, you are plunged into another. If, in your wanderings, you climb to the top of a tall tower, you can look down over the streets you have been lost in, and suddenly everything begins to make sense. You see where you should have turned one way but went another; you realize that the little shop you walked past, with the cat in the window, was only yards away from the garden in the next street, which you found hours later. And when you get back down into the maze you find your way easily. Now you know your way about.

In this book we shall find ourselves discussing the nature of morality, when we set out to decide whether it is always wrong to kill an innocent person; we shall end up talking about what it is for a theory to be scientific when we started out wondering about the claims of astrologers. And when this happens, I think it will help to bear in mind this image of being lost in an old city. When we move to these abstract questions, apparently remote from the practical concerns we started with, what we are doing is like climbing up that tower. From up there we can see our way around the problems. So that when we get back down into the city, back to the concrete problems that started us out, we should find it easier to get around.

People are normally introduced to philosophy by one of two routes. The first is through reading the more accessible of the great historical texts of philosophy—Plato's dialogues, for example, or Descartes' *Meditations*. The second is by examining some central philosophical question: "What is knowledge?" say, or "Is morality objective?" In this book I shall be following this second route, but I shall discuss the views of some of the great philosophers on the central questions on the way. Still, it is important to keep in mind that I will always be trying to move toward a philosophical understanding of the problem I am looking at, rather than trying to give a historically accurate account of a past philosopher.

It is fashionable, at the moment, to stress the way that the central problems of philosophy change over time. People say that no one nowadays can really be concerned with all of the problems that worried Plato. There is some truth in this. There are things in Plato that it is hard to understand or get excited by: much of the theory in the *Symposium* about the nature of love, for example, is likely to seem to a modern reader hopelessly wrong. Fortunately, however, a good deal more in Plato is extremely interesting and relevant: his *Theaetetus*, which is a dramatic dialogue about the nature of knowledge, remains one of the great classics of philosophy, and I shall discuss it in Chapter 2.

So the reason we philosophers continue to read Plato and many other philosophers between his time and ours is not simple curiosity about the history of our subject. Rather, we find in the great works of the past clues to a deeper understanding of the philosophical questions that trouble us now. That's why mentioning Plato and Descartes isn't some kind of concession to the proponents of the historical route into philosophy. It isn't even just a concession to old habits in the teaching of philosophy. It is simply a reflection of the facts that make the historical route work.

My aim in this book is twofold, then: First, I would like anyone who reads it carefully to be able to go on to read contemporary philosophical discussions. Second, I would like such a reader to be able when he or she comes to read Plato, say, or Descartes, to see why their work remains an enduring contribution to our understanding of the central problems of philosophy. I shall always have in mind a beginning philosophy student who knows none of the technical language of philosophy but is, nevertheless, willing to think through difficult questions. There are bibliographical notes and some advice on further reading at the end of the book; and there is also an index, which gives in bold type the page number of the page where a term is introduced or defined. Finally, because I often need to refer you back or forward to a discussion of a related issue, I have numbered the sections of each chapter. So sometimes I'll refer to section 5 of chapter 3, for example, as 3.5. Together, these various tools—the notes, the index, the further reading, and the numbered sections—are meant to help you find your way around.

You will learn a lot of new words in the course of reading this book. Philosophy, like all scholarly disciplines, has its own technical terms. We use them because technical language allows you to keep track of important distinctions and to speak and write in ways that are somewhat more precise than our everyday talk. The important thing is to grasp the ideas these terms express and the distinctions they make and to see how these distinctions and ideas can be used in arguments that deepen our understanding. And one general rule to keep in mind was set out by the Greek philosopher Aristotle about twenty-five hundred years ago: he insisted that we should adopt the degree of precision appropriate to the subject matter. We could say, more generally, that distinctions are worth making only if they do some work in an argument or help us to see something we wouldn't otherwise see. The technical terms are tools for a purpose, not the point of the exercise. As far as possible, contemporary philosophers actually prefer to use what the English philosopher Bernard Williams once called "moderately plain speech." So while philosophy has a technical vocabulary, doing philosophy means more than knowing and throwing around those special terms.

The book is organized around eight central areas of the subject: mind, knowledge, language, science, morality, politics, law, and metaphysics. (Only the last of these, as you see, has a technical name. When we get to the chapter on metaphysics, I'll explain why it has to be there.) In the chapter on language I say something about logic; in the chapter on metaphysics I discuss the existence of God.

Now I'm going to start straight in with Mind and this may seem surprising. You might have supposed that a good question to answer at the beginning of an introductory philosophy book is: "What is philosophy?" But I think that is a mistake, and if we consider the same question about a different subject, I think you will see why.

So consider the question: "What is physics?" If you asked what physics was, you might well get the answer that it is the study of the physical world. In some ways this isn't a very helpful answer. One trouble is that if you take the answer broadly, then biology is a

branch of physics: living organisms are part of the physical world. But this just shows that not every part of the physical world gets studied in physics. Which aspects *are* the physical aspects? Well, if you knew *that*, and were thus able to rule out biological questions, you would already be well on the way to knowing what physics is.

Nevertheless, there is a reason why most of us don't find this answer just unhelpful. We learned some physics in high school, and so we already have lots of examples of physical experiments and problems to draw on. These examples allow us to understand what is meant by "the physical world": it consists of those aspects of the world that are like the ones we studied in high school physics. If we tell someone who has never done any physics that physics is the systematic study of the physical world, we should not be surprised if they find our answer rather unhelpful.

There is a lesson here for how we should begin to develop an understanding of what philosophy is. What it suggests is that rather than tackling the question head on, we should look at some examples of philosophical work. With these examples in mind it won't be so unhelpful to be given an answer like the one we got to "What is physics?" For if we end up by saying that philosophy is the study of philosophical problems, that won't be uninformative if we have an idea of what some of the major philosophical problems are. So I'm not going to start this book by telling you what I—or anyone else—think philosophy is. I'm going to start by *doing* some. Just as you are in a better position to understand what physics is when you have done some, so you will be better able to see how philosophy fits into our thought and our culture when you have a "feel" for how philosophers argue and what they argue about.

Before we start I need, finally, to introduce a couple of conventions that I'm going to use. I shall use quotation marks to do two different jobs. One job—exemplified in the last sentence of the previous paragraph—is to indicate that a word is being used in a nonstandard way. Philosophers call these "scare quotes." The other job is to allow me to refer to words, sentences and other expressions, as when I say that the word "word" has four letters. The sentence

A: There are nine letters in "most words."

is true. The sentence

B: There are nine letters in most words.

is false. ("False," for example, has only five letters!) And I've just exemplified one other convention. When I display a sentence or expression indented on a line by itself, I will not put it in quotes; the fact of displaying it in this way is an alternative convention for allowing me to refer to words and other linguistic expressions. If I put a letter at the start of the line, I'll use that letter as the name of the sentence later. So here, for example, I can say that A and B have very different meanings. In A, we say, I am **mentioning** the words "most words." In B, I am **using** them. This distinction between use and mention may seem obvious. But sometimes, in a complex argument, we may get into a muddle if we don't keep use and mention distinct. In chapter eight, for example, we'll discuss the existence of numbers. There it will be important to distinguish between asking whether the numeral (i.e., the word or symbol) "9" exists, and whether 9 itself exists. The answer to the first question is obviously Yes. But the answer to the second question is not nearly so simple.

If I were to follow this convention strictly, then, when I introduced a term (as I often will) by saying "I will call something **X**," I would have to put the "X" in quotes. But here the boldface type can do the job of the quotes—which is to show that I'm *mentioning* a term and not *using* it—so I won't usually bother. The convention is meant to help avoid confusion: it's not an obsession to be pursued for its own sake! (For the record, terms occur in boldface only at the point where I introduce or define them.)

I began this introduction by mentioning various questions that might lead you to philosophy in the first place; but perhaps you have never been bothered by any such questions. That is no reason to think that philosophy is not for you. Many people do, of course, live their lives without ever thinking systematically about philosophy. But I shall be arguing that many problems that trouble us in ordinary life—down in the city, rather than up in the tower—can only

be answered if we first ask the more fundamental questions that are the hallmark of philosophy. Doing philosophy, then, enlarges your capacity to think about the life you are leading and what matters in it. Socrates famously said that the unexamined life was not worth living. Philosophy is one way to enrich your ability to examine the assumptions and ambitions that guide your life.

CHAPTER 1

Mind

What is a mind?
Could we make a machine with a mind?
What is the relationship between minds and bodies?

1.1 Introduction

In countless movies, computers play a starring role. Some talk in synthesized voices; others write a stream of words on a screen. Some manage spaceships; others, the "brains" of robots, manage their own "bodies." People converse with them, are understood by them, exchange information and greetings with them. Much of this is still science fiction. But real computers advise lawyers on relevant cases, doctors on diagnoses, engineers on the state of atomic reactors. Both the fantasy and the fact would have astonished our grandparents. *Their* grandparents might have thought that this could only be achieved by magic. Yet most of us are getting used to it, taking the silicon age for granted.

Still, a suspicion remains. We human beings have always thought of ourselves as special. We all assume some contrast between the world of material things and the world of spiritual things. If the computer really is a "material mind," then not only must we rethink this distinction, but we have broken it with our own creations. We should be careful to avoid such an important conclusion until we have really thought it through. However natural it seems to take it for granted that computers can think and act, then, we shouldn't just assume it. In philosophy we often find that what we normally take for granted—the "commonsense" point of view—gets in the way of a proper understanding of the issues. So let's see if the way I spoke about computers in the first paragraph is accurate.

I said that they talk. But do they *really* talk in the sense that

people do? It isn't enough to say that they produce something that sounds like speech. Tape recorders do that, but they don't talk. When people talk they mean something by what they say. To mean something, they need to be able to understand sentences. Now I also said that computers understand what we say to them. But do they really? The sounds of our speech are turned into electrical impulses. The impulses pass through the circuits of the machine. And that causes the speech synthesizer to produce sounds. It may be very clever to design a machine that does this, but what evidence do we have that the machine understands?

Well, could a machine understand? There are two obvious responses to this question. The first response I'll call **mentalist**, for the sake of a label. It's the response you make if you think that understanding what people say involves having a mind. The mentalist says:

> Computers can't really understand anything. To understand they would have to have conscious minds. But we made them from silicon chips and we programmed them. We didn't give them conscious minds. So we know they don't have them.

At the other extreme is the response I'll call **behaviorist**. The behaviorist says:

> Naturally, everyone should agree that some computers don't understand. But there's no reason why a computer couldn't be made that does understand. If a machine responds in the same ways to speech as a person who understands speech, then we have just as much reason to say that the machine understands as we have to say that the person does. A machine that behaves in every way as if it understands is indistinguishable from a machine that understands. If it behaved in the right way, that would show that it had a mind.

It is clear why I call this response "behaviorist." For the behaviorist says that to understand is to *behave* as if you understand.

What we have here is a situation that is quite familiar in philosophy. There are two opposing views—mentalist and behaviorist, in this case—each of which seems to have something in its favor, but neither of which looks completely right. Each of these views has a

bit of common sense on its side. The mentalist relies on the common sense claim that machines can't think. The behaviorist relies on the common sense claim that all we know about other people's minds we know from what they do. It looks as though common sense here isn't going to tell us if the mentalist or the behaviorist is right.

In fact, if you hold either of these views you can face difficult intellectual choices. Let's start with a problem you get into if you are a mentalist. Suppose the computer in question is in a robot, which, like androids in science fiction, looks exactly like a person. It's a very smart computer, so that its "body" responds exactly like a particular person: your mother, for example. For that reason I'll call the robot "M." Wouldn't you have as much reason for thinking that M had a mind as you have for thinking that your mother does? You might say, "Not if I know that it's got silicon chips in its head." But did you ever check that your mother has got brain tissue in her head? You didn't, of course, because it wouldn't prove anything if you did. Your belief that your mother has a mind is based on what she says and does. What's in her head may be an interesting question, the behaviorist will say, but it isn't relevant to deciding whether she has thoughts. And if it doesn't matter what is in your mother's head, why should it matter what's in M's?

That's a major problem if you're a mentalist: how to explain why you wouldn't say an android had a mind, even if you had the same evidence that it had a mind as you have that your mother does. Surely it would be absurd to believe your mother has a mind on the basis of what she does and says, yet refuse to believe M has a mind *on the very same evidence*. If it's the evidence of what your mother does that entitles you to believe she has a mind (and not, say, an innate prejudice), then the very same evidence about something else would entitle you to believe that *it* had a mind. This is one line of thought that might lead you to behaviorism.

But if you decide to be a behaviorist, you have problems too. You and I both know, after all, since we both do have minds, what it is like to have a mind. So you and I both know there's a difference between us and a machine that behaves exactly like us but doesn't have any experiences. Unless M has experiences, it hasn't got a mind. The difference between having a mind and operating as if

you've got one seems as clear as the difference between being conscious and being unconscious.

The upshot is this: If you look at the question from the outside, comparing M with other people, behaviorism looks tempting. From the point of view of the evidence you have, M and your mother are the same. Looked at from the inside, however, there is all the difference in the world. You know you have a mind because you have conscious experiences, an "inner life." M may have experiences, for all we know. But if it doesn't, no amount of faking is going to make it true that it has a mind.

We started with a familiar fact: computers are everywhere and they're getting smarter. It looks as though there will soon be intelligent machines, machines that will understand what we say to them. But when we look a little closer, things are not so simple. On the one hand, there is reason to doubt that behaving like a person with a mind and having a mind are the same thing. On the other, once we start asking what and how we know about the minds of other people, it seems that our conviction that people have minds is no better based than the belief that there could be understanding computers. We call someone who asks philosophical questions about what and how we know an **epistemologist**. And if we ask how we know about the minds of other people it seems plain that it is from what they say and do. We simply have no direct way of knowing what—if anything—is going on in other people's minds. But then, if what people say and do is what shows us they have minds, a machine that says and does the same things shows us that it has a mind also. From the epistemologist's point of view, other people's minds and the "minds" of computers are in the same boat.

When we look at the question from the inside, as we have seen, the picture looks different. Someone who looks from the inside we can call a **phenomenologist**. "Phenomenology" is the philosopher's word for reflecting on the nature of our conscious mental life. From the phenomenologist's point of view, M, and all machines, however good they are at behaving like people, may well turn out not to have minds.

From thinking about computers in science fiction we have found our way to the center of the maze of problems that philosophers call the **philosophy of mind** or **philosophical psychology**.