Praise for the previous edition:

"A really excellent introduction to philosophy does the following: meets the student at their level, then takes them up a notch, and approaches traditional topics in unique and interesting ways. This book does those things."
Fred Adams, *University of Delaware*

Writtten in an engaging lecture-style format, this 8th edition of *Core Questions in Philosophy* shows students how philosophy is best used to evaluate many different kinds of arguments and to construct sound theories. Well-known historical texts are discussed, not as a means to honor the dead or merely to describe what various philosophers have thought but to engage with, criticize, and even improve ideas from the past. In addition—because philosophy cannot function apart from its engagement with the wider society—traditional and contemporary philosophical problems are brought into dialogue with the physical, biological, and social sciences. Text boxes highlight key concepts, and review questions, discussion questions, and a glossary of terms are also included.

Core Questions in Philosophy has served as a premier introductory textbook for three decades, with updates to each new edition.

Key updates to this 8th edition include:

- A new chapter, "Probability and Bayes' Theorem"
- A new explanation of the concept of "soundness," as a useful tool in assessing arguments
- A clearer explanation, in the chapter on evolution, of the crucial biological idea that the similarities of different species provide evidence of their common ancestry
- A new discussion of evolutionary altruism in the chapter on psychological egoism
- A presentation of two interesting arguments from historically important Islamic and Confusian philosophers
- Improved clarity and updated material from philosophy and empirical research, throughout

Revisions to the online list of recommended resources include:

- Additional recommendations of supplementary readings, with the inclusion of more work from female philosophers
- New recommended videos and podcasts, all organized by their relevance to each chapter in the book

Elliott Sober is Hans Reichenbach Professor and William F. Vilas Research Professor in the Philosophy Department at the University of Wisconsin–Madison. His most recent book is *The Design Argument* (2018).

PHILOSOPHY

Cover image: © Fuat Oztas / Getty Images

Routledge
Taylor & Francis Group
www.routledge.com

eRESOURCES
www.routledge.com/cw/sober

Routledge titles are available as eBook editions in a range of digital formats

an **informa** business
ISBN 978-0-367-46498-1

CORE QUESTIONS IN PHILOSOPHY
Eighth Edition
ELLIOTT SOBER
Routledge

CORE QUESTIONS IN PHILOSOPHY

ELLIOTT SOBER

Eighth Edition

eRESOURCES
ROUTLEDGE

Core Questions in Philosophy

Written in an engaging lecture-style format, this 8th edition of *Core Questions in Philosophy* shows students how philosophy is best used to evaluate many different kinds of arguments and to construct sound theories. Well-known historical texts are discussed, not as a means to honor the dead or merely to describe what various philosophers have thought but to engage with, criticize, and even improve ideas from the past. In addition—because philosophy cannot function apart from its engagement with the wider society—traditional and contemporary philosophical problems are brought into dialogue with the physical, biological, and social sciences. Text boxes highlight key concepts, and review questions, discussion questions, and a glossary of terms are also included.

Core Questions in Philosophy has served as a premier introductory textbook for three decades, with updates to each new edition.

Key updates to this 8th edition include:

- A new chapter, "Probability and Bayes's Theorem"
- A new explanation of the concept of "soundness," as a useful tool in assessing arguments
- A clearer explanation, in the chapter on evolution, of the crucial biological idea that the similarities of different species provide evidence of their common ancestry
- A new discussion of evolutionary altruism in the chapter on psychological egoism
- A presentation of two interesting arguments from historically important Islamic and Confucian philosophers
- Improved clarity and updated material from philosophy and empirical research, throughout

Revisions to the online list of recommended resources include:

- Additional recommendations of supplementary readings, with the inclusion of more work from female philosophers
- New recommended videos and podcasts, all organized by their relevance to each chapter in the book

Elliott Sober is Hans Reichenbach Professor and William F. Vilas Research Professor in the Philosophy Department at the University of Wisconsin–Madison. His most recent book is *The Design Argument* (2018).

Praise for the previous edition:

"A really excellent introduction to philosophy does the following: meets the student at their level, then takes them up a notch, and approaches traditional topics in unique and interesting ways. This book does those things."

Fred Adams, *University of Delaware*

Core Questions in Philosophy

Eighth Edition

Elliott Sober

Routledge
Taylor & Francis Group
NEW YORK AND LONDON

Eighth edition published 2021
by Routledge
605 Third Avenue, New York, NY 10158

and by Routledge
2 Park Square, Milton Park, Abingdon, Oxon, OX14 4RN

Routledge is an imprint of the Taylor & Francis Group, an informa business

First edition published by Willan 2008
Seventh edition published by Routledge 2020

Library of Congress Cataloging-in-Publication Data
A catalog record for this book has been requested

ISBN: 978-0-367-46628-2 (hbk)
ISBN: 978-0-367-46498-1 (pbk)
ISBN: 978-1-003-03004-1 (ebk)

Typeset in Bembo
by Apex CoVantage, LLC

Access the eResource: www.routledge.com/cw/sober

For Aaron

Contents

Boxes

Preface

The philosophical problems investigated in this book concern fundamental facts about our place in the universe. Many of us were brought up to believe that God exists, that there is a real difference between right and wrong, that we can freely choose what sort of lives to lead, and that it is possible for us to gain knowledge of the world we inhabit. A major goal of philosophy is to discover whether these opinions can be rationally defended or are just comfortable illusions.

Core Questions in Philosophy emphasizes the idea that philosophy is a subject devoted to *evaluating arguments* and *constructing theories*. This is not the same as describing the history of what various philosophers have thought. Although I discuss historical texts, I do so because they are rich sources of ideas pertinent to answering philosophical questions. The point is not to say solemn and respectful words about worthy figures now dead, but to engage them in dialogue—to grapple with the theories they have proposed, to criticize these theories, even to improve upon them.

Besides proposing answers to philosophical questions, I also try to make clear which questions I have *not* answered. I hope that the reader will approach what I say in the way I have approached the philosophical texts I discuss. This is a book to argue with, to dissect. It isn't my goal to have the reader accept without question the conclusions I reach.

The chapters are intended to flow together, so that the main areas covered—philosophy of religion, epistemology, philosophy of mind, and ethics—are connected to each other to make a coherent whole. The chapters I wrote are intended to be *launching pads* from which readers can pursue issues on their own. I believe students are best able to think about philosophy if they are first provided with some basic tools and concepts. It is the purpose of the chapters to provide these *core ideas*.

Each chapter is followed by review questions and problems for further thought. These should help readers to consolidate their understanding of what I have said, and to think creatively about related problems. The chapters often contain material in boxes; these boxes provide, in a nutshell, a restatement of an important idea or a brief discussion of a related matter that may interest the reader. A list of the boxes immediately follows the table of contents. Each chapter of the book includes suggestions for further investigation. There is also a glossary at the end of the book that provides simple definitions of the main concepts used.

Besides discussing a number of traditional topics, this book also takes up some contemporary theories and problems, both from philosophy and from other disciplines. Creationism and evolutionary theory are hotly debated now. The issues they raise are continuous with a tradition of argument in philosophy of religion that goes back (at least) to Aquinas, Hume, and Paley. The relation of mind and body is a philosophical problem of long standing, but the ideas of Freud and Skinner get a hearing along with those of Descartes. In ethics, there has long been a debate as to whether ethical truths are discovered or created.

Plato and Sartre are separated by more than two thousand years, but both speak to this issue. The problem of free will raises the question of whether every event is caused. Here, the contribution of modern physics must be brought into contact with a perennial problem of philosophy. Philosophy isn't the same as biology, psychology, or physics, but the problems of philosophy cannot be isolated from the sciences. One aim of this book is to connect philosophical problems with ideas derived from a wider culture.

The etymology of the word *philosopher* is *lover of wisdom*. This doesn't guarantee that all philosophers are wise, nor even that each individual philosopher is devoted to the attainment of wisdom. Philosophers *should* strive for wisdom; whether they do so, and whether they attain it, are separate questions. Wisdom involves understanding—seeing how things fit together. When the pieces of a puzzle are fitted together, one attains a sense of wholeness. Current philosophy is embedded in a historical tradition of philosophical discourse. It also is connected to problems in the sciences, the other humanities, and the arts. This book aims to give the reader a sense of these multiple connections.

Elliott Sober
University of Wisconsin–Madison

Acknowledgments

My debts to colleagues in philosophy at University of Wisconsin, Madison, are enormous. A fixed point in my work week has been discussions of the ideas and techniques that go into presenting central problems of philosophy to new students. My philosophical outlook, as well as the view I have of teaching, have been shaped by these conversations.

It is a pleasure to acknowledge the help I've received from Michael Byrd, Claudia Card, Fred Dretske, Ellery Eells, Berent Enç, Malcolm Forster, Martha Gibson, Paula Gottlieb, Daniel Hausman, Andy Levine, Steve Nadler, Terry Penner, Mark Singer, Dennis Stampe, Daniel Wikler, and Keith Yandell. They were generous enough to suffer my trespasses onto philosophical terrain that belonged more to them than to me. Some read parts of this book and gave me comments; others listened patiently while I tried out what I thought was a new angle.

The first seven editions of *Core Questions in Philosophy* elicited a steady stream of correspondence and phone calls from teachers of philosophy and their students. These took a variety of forms; there was praise and blame, suggestions on how to do better, and even a few not-so-gentle suggestions that I should turn my attention to other projects. On the whole, though, I was happy with what I heard, though this didn't mean that I felt that I should leave the book unchanged. I thank everyone who took the trouble to let me know what they thought. Most (but not all) will find evidence that I listened to what they said in the way this edition differs from the ones before.

Deserving of special mention are Richard Behling, Keith Butler, John Carpenter, William R. Carter, Paul Christopher, Hayley Clatterbuck, Robert Cummins, Katie Deaven, Stewart Eskew, Doug Frame, Phil Gasper, Ronald Glass, Richard Hanley, Casey Hart, John Hines, Burton Hurdle, Paul Kelly, Charles Kielkopf, John Koolage, Matthew Maxwell, Gregory Mougin, Bradley Monton, Margaret Moore, William Russell Payne, Marcelo Pimentel, Howard Prospersel, David Ring, Roy Sorensen, Reuben Stern, Naftali Weinberger, Whilhelm S. Wurzer, Stephen Wykstra, Joel Velasco, and Shimin Zhao. The help they provided was extremely valuable. I also must thank the anonymous readers of previous editions of this textbook for their valuable suggestions. I also thank Sera Schwartz for her excellent work expanding the review questions found at the end of each chapter and updating the recommendations for supplementary readings, video, and audio, found on the eResource.

Writing an introduction to philosophy is a challenge. The challenge is to reconstruct what a problem or idea would sound like to someone who hasn't studied the subject before. The project requires that one return to the beginning—to the fundamentals of the subject. I hope what I found by beginning again will be useful to those who are beginning for the first time.

Routledge eResource

Provided with this text is a Routledge eResource.

Please visit at: www.routledge.com/cw/sober

This eResource is a list of recommended resources, including:

- **Recommended primary source reading:** Primary sources relevant to each chapter. Most are available to read for free online.
- **Recommended supplementary reading:** Reading suggestions for each chapter, comprising books, journal articles, encyclopedias, and blogs. These readings allow students to gain further understanding of the topics discussed in each chapter. Many of the readings are available to read for free online.
- **Recommended listening:** Links to free-to-listen podcasts and other audio recordings for most chapters, supporting learning.
- **Recommended watching:** Links to free-to-watch videos for most chapters, supporting learning.

Part I

Introduction

Chapter 1

What Is Philosophy?

Chapter Outline

When asked "do you have a philosophy?" most people say "yes," but what do they mean? They usually have in mind a set of beliefs that they admit are difficult to prove are true, but that nonetheless are important to the way they think of themselves and the world they inhabit. Sometimes people describe their philosophies by saying what they think makes an action right or wrong. The statement "it's part of my philosophy that people should help each other" might be an example. A person's philosophy might include the fundamental ethical principles he or she believes. But people often have more than *ethics* in mind when they talk about their philosophies. A religious person might say that it is part of his or her philosophy that God exists; an atheist might say that it is part of his or her philosophy that there is no God and that there is no life after death. These propositions are important to the people who believe them. They describe what exists; philosophers would say that they are part of *metaphysics*, not ethics. Metaphysics is the part of philosophy that attempts to describe, in very general terms, what there is.

If everyday people think of their philosophies as the important beliefs they have that are difficult to prove, how does this idea of philosophy relate to how philosophers understand their own subject? Sometimes a term is used in ordinary talk in a way that differs dramatically from the way it is used by specialists. People sometimes say that tomatoes are vegetables, but a botanist will tell you that tomatoes are fruits. Every day people say they are concerned about "ecology," but biologists understand "ecology" in a very different way. Perhaps philosophers use the term "philosophy" in a way that departs fundamentally from what ordinary people mean when they say that they have a philosophy.

To gain a better purchase on what philosophy is, I'm going to discuss the question of what is distinctive about philosophy from two angles. First, I'll sketch some of the main philosophical problems that I'll examine in this book. That is, I'll describe some *examples* of philosophy. But giving examples doesn't really answer the question of what philosophy is. If you asked, "What is a mammal?" and I showed you a human being, a hippo, and a cat,

these examples might give you a *hint* about what a mammal is. However, citing examples isn't the same as saying what it is to be a mammal. That is why there will be a second stage to my discussion of what philosophy is. After giving some examples of philosophical problems, I'll present some theories about what philosophy is. I believe these theories have merit, though I admit none is entirely adequate.

Examples

The first philosophical problem we'll consider in this book is whether God exists. Some philosophers have constructed arguments that attempt to establish that God exists, others have tried to show there is no God, and still others think that the question can't really be answered. I'll evaluate some of the more influential arguments and try to see whether they work.

The second problem we will consider concerns knowledge. It is pretty clear that belief and knowledge are different. Long ago some people thought that the earth is flat. They *believed* this, but they didn't know it, since it isn't true. Of course, they *thought* they knew it, but that's different. It is also pretty clear that true belief isn't the same as knowledge. If you believe something for no reason at all, but happen to be right by accident, you have true belief but not knowledge. For example, think of a gullible gambler at a racetrack who believes for no good reason that the first horse in every race will win. Occasionally this person will be right—she will have a true belief. But it isn't plausible to say that she knew, on those races about which she turned out to be right, which horse would win. So having knowledge involves something more than having a true belief.

The philosophical problem about knowledge will split into two parts. First, there are the questions: What is knowledge? What makes knowledge different from true belief? Second, there is the question: Do human beings ever know anything? One philosophical position we will consider answers this last question in the negative. Sure, we have beliefs. And granted, some of our beliefs turn out to be true. Knowledge, however, we never have. We don't even know those things that we take to be most obvious. This position is called *philosophical skepticism*. We will consider arguments for skepticism and arguments that attempt to refute it.

The third philosophical subject that will be addressed in this book consists of a collection of topics from the philosophy of mind. The first of these is the so-called mind/body problem. You have a mind; you also have a brain. What is the relationship between these items? One possible answer is that they are identical. Although "mind" and "brain" are different words, they name the same thing, just like the names "Superman" and "Clark Kent." An alternative position in this area is called *dualism*; it says that the mind and the brain are different things. We will consider other theories that have been advanced about the mind/body problem as well.

Another topic from the philosophy of mind that we'll address concerns human freedom. Each of us has the personalities we have because we inherited a set of genes from our parents and then grew up in a sequence of environments. Genes plus environments make us the sorts of people we are. We didn't choose the genes we have, nor did we choose the environments we experienced in early life. These were thrust upon us from the outside. Each of us performs certain actions and abstains from performing others. This pattern of what we do and don't do results from the personalities we have. Can we be said to perform actions freely? Is it really in our control to perform some actions and abstain from others? Perhaps the fact that our actions are the results of factors outside our control (our genes and our early environment) shows that it is a mistake to say that we freely choose what we do.

Of course, we talk in everyday life about people doing things "of their own free will." We also think of ourselves as facing real choices, as exercising control over what we do. However, the philosophical problem of freedom asks whether this common way of thinking is really defensible. Maybe freedom is just an illusion. Perhaps we tell ourselves a fairy tale about our own freedom because we can't face the fact that we aren't free. The philosophical problem will be to see whether we can be free agents if our personalities are the results of factors outside our control.

The last problem area we will address is ethics. In everyday life, we frequently think that some actions are right and others are wrong. The philosophical problem about this familiar attitude divides into two parts. First, we'll consider whether there really are such things as ethical facts. Maybe talk about ethics, like talk about freedom, is just an elaborate illusion. Consider a parallel question about science. In every science, there are questions that are controversial. For example, physicists have different opinions about how the solar system began. But most of us think that there is something else to physics besides opinions. There are facts about what the world is really like.

Clashes of opinion occur in what I'll call the *subjective realm*. Here we find one human mind disagreeing with another. But facts about physics exist in the *objective realm*. Those facts exist independently of anybody's thinking about them. They are out there, and science aims to discover what they are. In science, there are both subjective opinions and objective facts—people have beliefs, but there also exists, independently of what anyone believes, a set of facts concerning the way the physical world really is. The question about ethics is whether both these realms (subjective and objective) exist in ethics, or only one of them does. We know that people have different ethical opinions. The question is whether, in addition to those opinions, there are ethical facts. In other words, does ethics parallel the description I've just given of science, or is there a fundamental difference here? The accompanying two-by-two table illustrates this question. *Ethical subjectivism* is the philosophical thesis that there are no ethical facts, only ethical opinions. According to this position, the claim that "murder is always wrong" and the claim that "murder is sometimes permissible" are *both* misguided—there are no facts about the ethics of murder for us to have opinions about. We'll consider arguments supporting and criticizing this position.

	Subjective Realm	*Objective Realm*
Science	Scientific opinions	Scientific facts
Ethics	Ethical opinions	Ethical facts?

The second question that arises in ethics is this: If there are ethical facts, what are they? Here we assume a positive answer to the first question and then press for more details. One theory we'll consider is *utilitarianism*, which says that the action you should perform in a given situation is the one that will produce the greatest happiness for the greatest number of individuals. This may sound like common sense, but I'll argue that there are some serious problems with this ethical theory.

Three Theories about What Philosophy Is

I've just described a menu of four central philosophical problems: God, knowledge, mind, and ethics. What makes them all *philosophical* problems? Instead of giving examples, can we

say something more general and complete about what distinguishes philosophy from other areas of inquiry? I'll offer three theories about what is characteristic of at least some philosophical problems.

Several of the problems just described involve *fundamental questions of justification*. There are many things that we believe without hesitation or reflection. These beliefs that are second nature to us are sometimes called "common sense." Common sense says that the sense experiences we have (via sight, hearing, touch, taste, and smell) provide each of us with knowledge of the world we inhabit. Common sense also says that people often act "of their own free will," and common sense holds that some actions are right while others are wrong. Philosophy examines the fundamental assumptions we make about ourselves and the world we inhabit and tries to determine whether those assumptions are rationally defensible.

Another characteristic of many philosophical questions is that they are very *general*; often they're more general than the questions investigated in specific sciences. Physicists have asked whether there are electrons; biologists have investigated whether genes exist; geologists have sought to find out whether the continents rest on movable plates. However, none of these sciences really bother with the question of why we should think that physical objects exist. The various sciences simply *assume* that there are things outside the mind; they then focus on more specific questions about what those things are like. In contrast, it is a characteristically philosophical question to ask why you should believe that there is anything at all outside your mind. The idea that your mind is the only thing that exists is called *solipsism*. Philosophers have addressed the question of whether solipsism is true. This is a far more general question than the question of whether electrons, genes, or continental plates exist.

The third view of what philosophy is says that philosophy is the enterprise of *clarifying concepts*. Consider some characteristic philosophical questions: What is knowledge? What is freedom? What is justice? Each of these concepts applies to some things but not to others. What do the things falling under a concept have in common, and how do they differ from the things to which the concept does not apply?

We must be careful here, since many questions that aren't especially philosophical sound like the examples just given. Consider some characteristic scientific questions: What is photosynthesis? What is acidity? What is an electron? How does the first batch of questions differ from these? One difference between these questions concerns the ways in which *reason* and *observation* help answer them. You probably are aware that philosophy courses don't include laboratory sections. Philosophers usually don't perform experiments as part of their inquiries. Yet, in many sciences (though not in all), laboratory observation is central. This doesn't mean that observation plays no role in philosophy. Many of the philosophical arguments we will consider begin by making an observation. For example, in Chapter 5, I'll consider an argument for the existence of God that begins with the following assertion: Organisms are complicated things that are remarkably well adapted to the environments they inhabit. The thing to notice here is that this fact is something we know by observation. Philosophers, as well as scientists, rely on observations.

Nonetheless, there is something distinctive about how observations figure in a philosophical inquiry. Usually the observations that are used in a philosophical theory are familiar and obvious to everyone. A philosopher will try to show by reasoning that those observations lead to some rather surprising conclusions. That is, although philosophy involves both observation and reasoning, it is the latter that in some sense does more of the work. As you will see in what follows, philosophical disputes often involve disagreements about reasoning; rarely are such disputes decidable by making an observation.

Each of these ways of understanding what philosophy is should be taken with a grain of salt (or perhaps with two). I think there is something to be said for each, even though each is somewhat simplified and distorting.

The Nature of Philosophy Has Changed Historically

One thing that makes it difficult to define "what philosophy is" is that the subject has been around at least since the ancient Greeks and has changed a great deal. There are many problems that are just as central to philosophy now as they were to the ancient Greeks, but there are other problems that have broken away from philosophy and now are thought of as purely scientific.

For example, ancient Greek philosophers discussed what the basic constituents of physical things are. Thales (who lived around 580 B.C.E.) thought that everything is made of water; many other theories were discussed as well. Now such questions are thought to be part of physics, not philosophy. Similarly, until the end of the nineteenth century, universities put philosophy and psychology together in the same academic department. It is only recently that the two subjects have been thought of as separate. Scientists in the seventeenth century—for example, Isaac Newton—used the term "natural philosophy" to refer to what we now think of as science. The term "scientist" was invented in the nineteenth century by the British philosopher William Whewell. The idea that philosophy and science are separate subjects may seem clear to us now, but the separation we now find natural was not so obvious in the past. Many of the problems that we now regard as philosophical are problems that have not broken away from philosophy and found their way into the sciences. Perhaps there are problems now taken to be philosophical that future generations won't regard as such. The shifting historical nature of what counts as philosophy makes it difficult to say anything very precise about what that subject is.

Philosophical Method

Having tried to say something about what philosophy is, I now want to say something about what philosophy is *not* (at least not in this book). You may have the impression that doing philosophy involves lying under a tree, staring up at the sky, and making deep and mysterious pronouncements off the top of your head that sound very important but that are hard to make sense of when you try to think about them clearly. I'll call this the *mystical guru model* of philosophy. Your experience reading this book won't correspond to this impression.

There is, however, another experience you've probably had that comes closer. If you took a high school geometry course, you'll remember proving theorems from axioms. If your geometry course was like the one I had, the axioms were given to you with very little explanation of why you should believe them. Maybe they looked pretty obvious to you, and so you didn't wonder very much about their plausibility. Anyhow, the main task was to use the axioms to prove theorems. You started with the axioms as assumptions and then showed that if they are true, other statements must be true as well.

Philosophers tend to talk about "arguments" rather than "proofs." The goal is to try to reach answers to important philosophical questions by reasoning correctly from assumptions that are plausible. For example, in Chapter 4, I'll examine some attempts to prove that God exists. The idea here is to start with assumptions that practically anybody would grant are true and then show that these assumptions lead to the conclusion that there is a God. This resembles what you may have done in geometry: Starting with simple and supposedly

obvious assumptions, you were able to establish something less obvious and more complex—for example, that the sum of the angles of a triangle equals two right angles (180°).

Sometimes the philosophical questions we'll consider will strike you as difficult, deep, even mysterious. I won't shy away from such questions. I'll try, however, to address them with clarity and precision. The goal is to take hard questions and deal with them clearly, which, I emphasize, should never involve trying to pull the wool over someone's eyes by making deep-sounding pronouncements that mean who-knows-what.

Summary

I began this chapter by describing how every day people use the term "philosophy." In fact, their usage is not so distant from what philosophers mean by the term. Philosophy *does* address the most fundamental beliefs we have about ourselves and the world we inhabit. Precisely because these assumptions are so central to the way we think and act, it is difficult to step back for a moment from these assumptions and examine them critically. The French have an expression: "the most difficult thing for a fish to see is water." Some assumptions are so natural and seemingly obvious that it is hard to see that we are making assumptions at all. Philosophy is the effort to help us identify these assumptions and evaluate them. Each of us *does* have a philosophy. What divides some people from others is their willingness to ask probing questions about what they believe and why. This is what philosophy as a discipline tries to add to the philosophies that each of us carries with us through our lives.

Review Questions

1 What is the difference between objective and subjective?
2 If you want to say what philosophy is, why isn't it enough to list some examples of philosophical problems?
3 What is the difference, if any, between "having a philosophy" and "doing philosophy"?
4 How do disagreements in reasoning differ from disagreements about observations? How do you think such disagreements are typically resolved?

A Problem for Further Thought

Which of the ideas presented here about what philosophy is also apply to mathematics? Which do not?

Recommended Readings, Video, and Audio

Visit the eResource at www.routledge.com/cw/sober for suggestions of readings, video, and audio, for this chapter.

Chapter 2

Deductive Arguments

Chapter Outline

Philosophy involves constructing and evaluating arguments. In this respect, philosophy is no different from any other rational activity—mathematicians do this, as do economists, physicists, and people in everyday life. The distinctive thing about philosophy isn't that philosophers construct and evaluate arguments; what is distinctive is the kinds of questions those arguments aim to answer. In the previous chapter, I talked about what makes a question philosophical. The goal in this chapter is to develop some techniques that can be used to tell whether an argument is good or bad.

Arguments

An argument divides into two parts: the premises and the conclusion. The premises and the conclusion are statements; each is expressed by a declarative sentence. Each is either true or false. When people argue that a given statement is true, they try to provide reasons for thinking this. The reasons are the premises of their argument; premises are assumptions. The statement to be established is the argument's conclusion.

In high school geometry, you talked about axioms and theorems. Axioms are assumptions (premises); the theorem (the conclusion) is what is supposed to follow from those assumptions. In geometry you may have spent little or no time asking whether the axioms are true. Not so for the philosophical arguments I discuss in this book. We'll want to see whether the premises are plausible. We'll also want to see whether the premises, if they were true, would provide a reason for thinking the conclusion is true as well. I'll pose these two questions again and again.

Good Arguments

I now want to talk about different kinds of "good arguments." What does "good" mean? A good argument is *rationally persuasive*; it gives you a substantial reason to think the conclusion is true. Advertisers and politicians sometimes use arguments that trick people into believing what they say. These arguments sometimes persuade people, but they don't always provide *good* reasons.

A good argument should have true premises; if the premises are false, how could they give you good reasons to believe the conclusion? But more is required than this. In the following argument, the premise is true, but it doesn't provide you a good reason to think that the conclusion is true:

Grass is green.

Roses are red.

What is wrong here is that the premises are irrelevant to the conclusion. A good argument should contain true premises, but it should also cite premises that are related in the right way to the conclusion. The truth of the premises should give you a reason to think that the conclusion is true. The three types of "good argument" that I'll now describe differ in what relationship their premises and conclusions have to each other.

Good arguments can be divided into two categories, and one of those categories can be divided into two more:

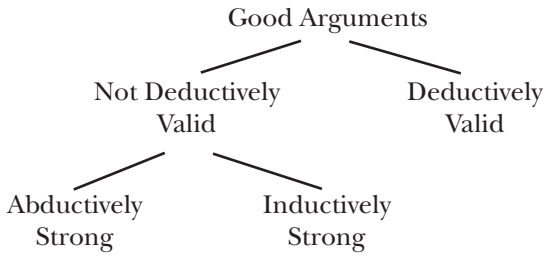

I'll treat the three categories (deductively valid, inductively strong, and abductively strong) as *mutually exclusive*. If an argument belongs to one category, it can't belong to any of the others. At the end of Chapter 3, I'll modify this classification slightly.

You may have heard some of this terminology before. Deduction is what you do in a mathematical proof. Induction involves sampling from a population to decide what its characteristics are. "Abduction" may be a less familiar term. It has nothing to do with kidnapping. The word was invented by the great nineteenth-century American philosopher Charles Sanders Peirce. Philosophers sometimes use the longer label "inference to the best explanation" to describe what Peirce meant by abduction.

I'll consider deduction in this chapter, induction and abduction in the next. The goal in each case is to describe some of the considerations that are relevant to deciding whether an argument is good or bad.

Deductive Validity Defined

The first type of good argument consists of ones that are "deductively valid." Here are two examples of this type of argument:

All fish swim.	All particles have mass.
All sharks are fish.	All electrons are particles.
All sharks swim.	All electrons have mass.

In these arguments, the premises are the statements above the horizontal line; the conclusion is the statement below. These arguments say that the premises are true and that, therefore, the conclusion also is true.

Here is what deductive validity means:

> *A deductively valid argument* is an argument that has the following property: *IF* its premises were true, its conclusion would have to be true.

I've capitalized the word *IF*. I'd print it in bright colors if I could because it is important not to forget this two-letter word. *A valid argument need not have true premises*. What is required is that the conclusion would have to be true *IF* the premises were true. Take a minute to look at these two arguments. Convince yourself that they are deductively valid.

"Validity" Is a Technical Term

What philosophers and logicians mean by "valid" doesn't have much in common with what we mean by "valid" in ordinary English. In everyday life, we say that a statement is "valid" if it is plausible or true. The technical use of the term that I just explained differs from ordinary usage in two ways. First, we never say that a *statement* or an *idea* is valid or invalid. Validity is a property of *arguments* and of arguments only. Second, an argument can be valid even if the statements it contains are wildly implausible. A valid argument can have false premises and a false conclusion.

Here is an example:

> All plants have minds.
> All ladders are plants.
> ---
> All ladders have minds.

Logical Form

What makes an argument deductively valid? The three example arguments described so far have different subject matters. The first is about fish, the second is about particles, and the third is about plants. Although they are about different things, they have the same structure. The structural property that they have in common is called their "logical form." Think of each argument as the result of substituting terms into the following skeleton:

> All *B*s are *C*s
> All *A*s are *B*s
> ---
> All *A*s are *C*s

This is the logical form of the three arguments given. You can think of *A, B*, and *C* as blanks into which terms may be substituted. Take a minute to see how the arguments just stated can be obtained from the above skeleton by substitution—by "filling in the blanks."

An argument is valid or invalid solely because of the logical form it has. The subject matter of the argument is irrelevant. Since the three example arguments have the same logical form, they are all valid or all invalid. They have the same logical form, so they are in the same boat. As already mentioned, they are valid. Indeed, each and every one of the millions of arguments you can construct by substitution into the above skeleton is valid as well.

Invalidity

The definition of validity tells you what a deductively *invalid* argument will be like. If there is even the smallest possibility that the conclusion could be false when the premises are true, then the argument is deductively invalid.

The ladder argument is valid, although all the statements it contains are false. Is the reverse situation possible? Can an argument be *invalid*, even though all the statements it contains are true? The answer is *yes*. Here's an example:

Emeralds are green.

Lemons are yellow.

The premise is true, and so is the conclusion. So why isn't the argument deductively valid? The definition of validity says that the premises in a valid argument must provide an *absolute guarantee* that the conclusion is true. But the fact that emeralds are green doesn't guarantee that lemons must be yellow. The color of lemons isn't entailed by the fact that emeralds are green. Validity concerns the *relationship* of premises to conclusion, not the question of whether the premises and the conclusion each happen to be true.

Sometimes it isn't so obvious that an argument is invalid. The above example is pretty blatant—the premise has nothing to do with the conclusion. But what do you think of the following argument? Is it valid or not?

If Jones stands in the heavy rain without an umbrella, then Jones will get wet.

Jones is wet.

Jones was standing in the heavy rain without an umbrella.

Imagine that all three of the statements in this argument are true. Imagine that Jones is now standing before you soaking wet and that Jones just came in from the rain.

Even if all the statements in this argument are true, this argument is still invalid. It is just like the argument about emeralds and lemons. Although the premises and the conclusion happen to be true, the premises don't *guarantee* that the conclusion must be true.

How can we see this more clearly? I said before that all arguments that have the same logical form are in the same boat. This means that if the argument about Jones is invalid, so is each and every argument that has the same form. Let's begin by isolating the argument's logical form. Here it is:

If *P*, then *Q*
Q

P

What do *P* and *Q* stand for in this argument skeleton? You can substitute any statement (declarative sentence) you please for these letters to obtain an argument with this logical

form. Notice that the letters in this skeleton differ in their function from the letters in the previous skeleton. There, *A, B*, and *C* were blanks into which terms denoting kinds of things ("fish," "electrons," etc.) can be substituted. Anyhow, we now have the logical form of the argument about Jones. If it is invalid, so are *all* arguments that have the same logical form. This means that if there is even one argument that has this logical form in which the premises are true and the conclusion false, then the argument form is invalid. This will mean that the initial argument about Jones is invalid as well.

Here is an argument that has the same logical form as the argument about Jones that settles the question:

> If Sam lives in Wisconsin, then Sam lives in the United States.
> Sam lives in the United States.
> ---
> Sam lives in Wisconsin.

The premises of this argument are true, but the conclusion, I assure you, is false. The Sam I'm talking about lives in Georgia.

Testing for Invalidity

Here's a strategy to use if you want to know whether an argument is invalid: First, ignore the argument's subject matter and isolate the logical form (the "skeleton") of the argument. Second, see if you can invent an argument that has this logical form in which the premises are true and the conclusion is false. If you can find even one rotten apple of this type, you are finished. If there is even one argument with this property, then *every* argument of that form is invalid.

When an argument has true premises and a false conclusion, it is quite obvious that the truth of the premises doesn't *guarantee* that the conclusion must be true. The premises can't be guaranteeing this, as the conclusion is false. This tells you something general. It tells you that each and every argument of this form will be such that the premises don't guarantee the truth of the conclusion.

So far, I've presented some examples of arguments. I've explained that a valid argument needn't have true statements in it and that an argument composed solely of true statements needn't be valid. This should make you wonder whether there is any connection at all between the question of whether an argument is valid and the question of whether the premises and conclusion are true.

There *is* a connection. It is illustrated by the following table. If an argument is valid, it can exhibit three of the four following combinations in which the premises are either all true or not all true and the conclusion is either true or false:

		Premises	
		All true	Not all true
Conclusion	True	Possible	Possible
	False	Impossible	Possible

This table indicates that a valid argument can't have true premises and a false conclusion. However, the fact that an argument is valid leaves open which of the other three cells in the table the argument occupies.

What can be said of an *in*valid argument? If an argument is invalid, are any of the four combinations impossible? I leave this to you to figure out by consulting the definition of validity.

When you find an invalid argument, you may want to ask if the argument can be repaired. Is there anything that can be done to an invalid argument to turn it into an argument that is valid? There is. By adding premises, you can always turn a deductively invalid argument into a valid one. Consider the following argument:

Smith lives in the United States.

Smith lives in Wisconsin.

This is invalid, but it can be made valid by adding a premise:

Smith lives in the United States.
Everyone who lives in the United States lives in Wisconsin.

Smith lives in Wisconsin.

Notice that the conclusion now follows from the premises. The trouble is that the second premise is false.

In the preceding pair of arguments, fixing the defect of invalidity just substitutes one problem for another; instead of having to criticize an argument for being invalid, you now have to criticize an argument for having a premise that isn't true. The following argument pair is different. Here you can repair the defect of invalidity and obtain a perfectly fine argument. Notice first that the following argument is not deductively valid:

Smith lives in Wisconsin.

Smith lives in the United States.

The argument can be repaired, however, by adding a premise:

Smith lives in Wisconsin.
Everyone who lives in Wisconsin lives in the United States.

Smith lives in the United States.

This argument is valid and has true premises as well. You can see from these two pairs of arguments that invalidity is easy to fix. Just add premises. What is harder is to add premises that not only make the argument valid, but that are true as well.

This idea will come up repeatedly when I discuss various philosophical arguments. I will sometimes claim an argument is invalid. When this happens, you should ask yourself whether the argument can be repaired. Often the price of making the argument valid (by adding a premise) is that you have to supply a new premise that you think is false. In making this addition, you are trading one defect (invalidity) for another (false premises). If you can't

repair the argument so that it is both valid and has all true premises, then you should consider the possibility that there is something fundamentally flawed about the whole line of argument. On the other hand, sometimes an invalid argument can be replaced by a valid one merely by supplying a true premise that maybe you neglected to mention because it is so obvious. In this case, the defect in the original argument isn't fundamental.

So far I've emphasized two questions that we will want to ask about arguments:

1 Is the argument deductively valid?
2 Are all the premises true?

If the answer to both questions is *yes*, the conclusion of the argument must be true. The technical term in logic for an argument that is both deductively valid and has true premises is soundness.

Arguments are tools. We use them to do things. When the goal is rational persuasion, a good argument will provide a good reason to think that the conclusion is true. If an argument is deductively valid and has true premises, is that sufficient to make the argument good? To see why validity and true premises aren't enough, consider the following argument:

Lemons are yellow.

Lemons are yellow.

Here the conclusion merely repeats what the premise asserts. This argument is valid and the premise is true. But there is something defective about this argument. What is it?

1 Conditionals

If/then statements are called conditionals. Conditional statements have other statements as components. For example, the statement "If pigs fly, then grass is green" is a statement of the form "If *P*, then *Q*," where *P* and *Q* are themselves statements.

In the statement "If *P*, then *Q*," *P* is called the *antecedent* and *Q* is called the *consequent*. A conditional doesn't say that its antecedent is true; the statement "If Joe drinks arsenic, then Joe will die" does not say that Joe drinks arsenic. And a conditional doesn't say that its consequent is true; "If there is a nuclear war, then Washington will be attacked" doesn't say that Washington will be attacked.

Conditionals can be rewritten without changing what they say. Consider the statement "If you live in Wisconsin, then you live in the United States." This is equivalent in meaning to "If you don't live in the United States, then you don't live in Wisconsin." The conditional "If *P*, then *Q*" is equivalent to "If not-*Q*, then not-*P*," no matter what *P* and *Q* happen to be. Here's a piece of terminology: The statement "If not-*Q*, then not-*P*" is the *contrapositive* of the conditional "If *P*, then *Q*." A conditional and its contrapositive are equivalent.

Consider the following two conditionals: "If *P*, then *Q*" and "If *Q*, then *P*." Are they equivalent? That is, do they mean the same thing? The answer is *no*. "If you live in Wisconsin, then you live in the United States" is true, but "If you live in the United States, then you live in Wisconsin" is false. These two if/then statements can't mean the same thing, because one is true while the other is false. "If *Q*, then *P*" is called the *converse* of the conditional "If *P*, then *Q*." A conditional and its converse are not equivalent.

Circularity, or Begging the Question

The previous argument is *circular*; it *begs the question*. Suppose you didn't already have an opinion as to whether lemons are yellow. The above argument wouldn't help you resolve your uncertainty. The argument would be useless in this regard.

Good arguments are tools that help answer questions about whether their conclusions are true. A good argument should give you a reason to accept the conclusion if you don't already believe the conclusion is true. So besides checking to see if an argument is deductively valid and has true premises, you should also see if the argument begs the question.

You'll notice in what I've just said that I am using the expression "begging the question" to name a defect in an argument. Unfortunately, the phrase is often used now to simply mean that a question is being asked. You'll hear this usage on the evening news—"the recession has not improved. This begs the question of whether government action can make things better." As with the term "validity," the term "begging the question" is used by philosophers with a special, technical meaning, one that doesn't coincide with ordinary usage.

Truth

One other idea needs clearing up before I leave the topic of deductive validity. You'll notice that the definition of validity makes use of the concept of truth. What is truth?

There are deep philosophical questions here, most of which I'll skirt. My goal is to describe the concept of truth I use in this book. It is beyond the scope of this book to defend this choice or to fully develop its implications. To begin with, whether a statement is true is an entirely different question from whether you or anybody happens to believe it. Whether someone believes the statement "The Rocky Mountains are in North America" is a psychological question. If beings with minds had never populated the earth, no one would have thought about the location of this mountain range. But this doesn't affect the question of whether the statement is true. There can be truths that no one believes. Symmetrically, there can be propositions that everyone thinks are true, but that aren't. There can be beliefs that aren't true.

2 Begging the Question

To understand what makes an argument question-begging, it is useful to examine some examples.

Suppose you were trying to convince someone that God exists. The argument you give for thinking that this is true is that the Bible says that there is a God. Would this argument convince someone who didn't already believe that there is a God? Probably not. Anyone who doubts that there is a God probably doesn't think that everything the Bible says is true.

Here's a second example. Someone is very suspicious about the reliability of consumer magazines. You try to convince him that *Consumer Reports* is reliable by pointing out that *Consumer Reports* ranks itself very highly in an article evaluating the reliability of consumer magazines. Probably your argument will fail to convince.

In these two examples, identify the premises and conclusion in each argument. Then describe what it is about the argument that makes it question-begging.

When I say that a certain sentence has the property of being true, what am I saying? For example, when I say that the English sentence "The Rocky Mountains are in North America" is true, am I attributing some mysterious property to the sentence? Not really.

All I'm saying is that the world is the way the sentence says it is. When I say that the sentence is true, all I'm saying is that the Rocky Mountains are in North America. So in a way, the concept of truth is often "redundant." Sometimes when I use the concept of truth, I could say the same thing without using that concept.

In high school English, your teacher might have told you to avoid redundancy. If you hand in an essay containing the sentence "Oscar is an unmarried bachelor," the essay might come back with "unmarried" crossed out and the marginal comment "avoid redundancy." The word "unmarried" is redundant because "Oscar is an unmarried bachelor" means exactly the same thing as "Oscar is a bachelor." Adding the word "unmarried" is to spill useless ink. The Redundancy Theory of Truth claims that the word "true" is redundant in just this sense. "It is true that the Rockies are in North America" says exactly what the sentence "The Rockies are in North America" asserts. This helps show why truth isn't a mysterious property. If you believe a statement *P*, you also believe that *P* is true. So, if you have any beliefs about the world at all, you should be quite comfortable applying the concept of truth to those beliefs.

"True for Me"

You'll see from these remarks that the expression "It is true for me" can be dangerously misleading. Sometimes saying that a statement is true "for you" just means that you believe it. If that is what you want to say, just use the word "belief" and leave truth out of it. However, there is a more controversial idea that might be involved here. Sometimes people use the expression "true for me" to express the idea that each of us makes our own reality and that the beliefs we have constitute that reality. I'll assume this is a mistake. My concept of truth assumes a fundamental division between the way things really are and the way they may seem to be to this or that individual. This is what I meant in Chapter 1 by distinguishing the objective realm and the subjective realm.

Wishful Thinking

Closely related to this distinction between objective and subjective is a piece of advice: *We should avoid wishful thinking*. Most of the things we believe aren't made true by our believing them. That the Rockies are in North America is a fact that is independent of our thought and language. We don't bring this geographic fact into being by thinking or talking in the way we do.

Self-Fulfilling Prophesies

In saying this, I'm not denying that the thoughts we have often affect the world outside the mind. If I think to myself, "I can't hit a baseball," this may have the effect that I do badly in the batter's box; here my believing something has the effect that the belief is made true. This is the idea of a "self-fulfilling prophesy." Notice how this causal chain works:

Thought		**Action**		**Truth**
I believe that I won't hit the baseball.	⟶	I swing too high.	⟶	I don't hit the baseball.

My believing a proposition causes an action, which has the effect of making the proposition true.

I have no problem with the idea that various statements may be caused to be true by individuals thinking thoughts to themselves. What I deny is that the mere act of thinking, unconnected with action or some other causal pathway, can make statements true in the world outside the mind. I'm rejecting the idea that the world is arranged so that it spontaneously conforms to the ideas we happen to entertain.

Later in this book, I'll investigate whether there are any exceptions to this principle that says that we should avoid wishful thinking. Maybe there are some statements that become true just because we think they are. Here are some philosophical claims we'll consider:

- Mathematical statements and definitions are made true by our regarding them as such; for example, "2 + 3 = 5" is true just because we choose to define our terminology ("2," "+," etc.) in the way we do (Chapter 4).
- Some statements about the contents of your own mind (for example, "You are in pain") are made true just by your believing they are true (Chapter 13).
- Ethical statements are true just because God, society, or some individual agent thinks they are (Chapter 32).

I'm mentioning these philosophical claims here without tipping my hand as to whether I think any of them is plausible. If any of them were correct, they would be exceptions to the pattern I've just described. For the moment, though, I'm merely noting that belief and truth are generally very separate questions.

Review Questions

1 When is a statement or idea valid? (a trick question)
2 Define what it means to say that an argument is deductively valid.
3 Invent an example of a valid argument that has false premises and a true conclusion. Invent an example of an invalid argument that has true premises and a true conclusion.
4 Can a statement be a premise in one argument and a conclusion in another? If you think so, give an example.
5 Which of the following argument forms is valid? Which is invalid? For each of the invalid ones, construct an example of an argument with that form in which the premises are true and the conclusion false:

(a)
If *P*, then *Q*
P
———————
Q

(b)
If *P*, then *Q*
Q
———————
P

(c)
If *P*, then *Q*
Not-*P*
———————
Not-*Q*

(d)
If *P*, then *Q*
Not-*Q*
———————
Not-*P*

For the argument forms you think are fallacious, invent names for these fallacies by using the vocabulary about conditionals presented in the box on page 18.

6 A sign on a store says, "No shoes, no service." Does this mean that if you wear shoes, then you will be served? What conditional statement does the sign express?

7 What does it mean to say that an argument is "circular," that it "begs the question"? Construct an example of an argument of this type different from the ones presented in this chapter.
8 What does it mean to say that truth is objective, not subjective?
9 What is the logical difference between the statements "I believe that *p*" and "*p* is true"? (Hint: Can you imagine a situation in which the latter, but not the former, is true? A situation in which the latter, but not the former, is false?)
10 Can an argument be "partially" or "somewhat" valid? Can a premise or conclusion be "partially" or "somewhat" true? Why or why not?
11 What does it mean to say that the premises of an argument "support" its conclusion?
12 Why do philosophers distinguish between considerations of "validity" and "truth"?
13 In this chapter I said that there are three types of good argument, and that one of them involves arguments that are deductively valid. Does that mean that all deductively valid arguments are good?

Problems for Further Thought

1 The Redundancy Theory of Truth may seem plausible as an account of what the following sentence means:

> It is true that the Rockies are in North America.

Does it work as well as an explanation of what the following sentence means?

> Some statements that are true have not been formulated yet.

2 Consider the following argument:

> I release an otherwise unsupported apple from my hand a few feet from the earth's surface.
>
> ---
>
> The apple falls to earth.

Is this argument deductively valid? What is the logical form of this argument?

3 (Here is a problem that was drawn to my attention by Richard Behling.) In this chapter, I said that each argument has a *single* logical form. This is the skeleton into which terms can be substituted to obtain the argument. What I said is an oversimplification. A given argument can be obtained from *many* logical forms. For example, consider the following argument:

> Fred lives in California.
> If Fred lives in California, then Fred lives in the United States.
>
> ---
>
> Fred lives in the United States.

This argument can be obtained from *both* of the following skeletons by substitution:

(a)
X
If *X*, then *Y*

Y

(b)
R
S

T

Argument form (a) is valid, but (b) is *in*valid. The argument about Fred is valid.

Here's the problem: Use the concept of logical form to define when an argument is valid, and when it is invalid, without falling into the trap of thinking that each argument has *exactly one* logical form.

4 In this chapter, I claimed that there are "objective truths." Do you agree? Construct an argument in which you try to demonstrate that such things exist, or that they do not.

Recommended Readings, Video, and Audio

Visit the eResource at www.routledge.com/cw/sober for suggestions of readings, video, and audio, for this chapter.

Chapter 3

Inductive and Abductive Arguments

Chapter Outline

In Chapter 2, I explained the idea of deductive validity. In a deductively valid argument, the premises provide an *absolute guarantee* that the conclusion is true: If the premises are true, there is no way in the world that the conclusion can be false.

Deductive Validity Is a Limitation

This feature of deductive arguments may sound like a virtue. It is a good thing when an argument provides this sort of strong guarantee. This virtue, however, can also represent a kind of limitation. Granted, a deductively valid argument that has true premises can't have a false conclusion; but it is also a property of such arguments that the conclusion can't say anything that wasn't already contained in the premises.

To see what this means, consider what you could validly deduce from the result of an opinion survey. Suppose you were interested in finding out what percentage of registered voters in a county are Democrats. You don't feel like contacting each of them and asking, so you open the phone book and make, let's say, one thousand telephone calls.

Suppose the result of your survey is that 60 percent of the people called say they are Democrats. What you want to know is the percentage of Democrats in the whole county. Could you construct a deductively valid argument here? Can you deduce that (approximately) 60 percent of the voters in the county are Democrats from a premise that describes the result of your survey? The answer is *no*, for two reasons. The fact that 60 percent of the people called *said* they are Democrats doesn't deductively guarantee that any of them are. And even if 60 percent of the people called are Democrats, you can't validly deduce from this that (approximately) 60 percent of the voters in the county are Democrats. That is, neither of the following arguments is deductively valid:

60 percent of the people called said they are Democrats.

---

60 percent of the people called are Democrats.

60 percent of the people called are Democrats.

Approximately 60 percent of the voters in the county are Democrats.

Why can't you validly deduce these conclusions? The reason is that in a deductively valid argument, it is impossible for the conclusion to be false if the premises are true. But it *is* possible that everyone you called in your survey lied. In addition, it is possible that the percentage of Democrats in the whole county is only 25 percent, even if nobody lied in your phone survey. In saying this, I'm not saying that the people you called actually lied, and I'm not saying that the real percentage in the whole county is only 25 percent. I'm just saying that these are *possible*, given the result of your phone calls. The result of your telephone survey doesn't absolutely rule out these possibilities; this means that you can't *deduce* the percentage of Democrats in the whole county from what the one thousand people said on the phone.

So the absolute guarantee that a deductively valid argument provides has this limitation: Insisting that an argument be deductively valid prohibits you from reaching conclusions that go beyond the information given in the premises.

It would make sense to insist that an argument be deductively valid if you wanted to avoid even the smallest risk of having a false conclusion with true premises. However, we are often willing to gamble. For example, we might think that the result of the phone survey does provide information about the composition of the county. We might think that the survey provides a pretty good reason for concluding that about 60 percent of the county voters are Democrats. However, the "good reason" isn't a deductively valid one.

Nondeductive Inference—A Weaker Guarantee

We have here a fundamental characteristic of nondeductive inference. Suppose we conclude that about 60 percent of the voters in the county are Democrats, based on the premise that 60 percent of the people called said they were Democrats. In this case, the premise doesn't provide an absolute guarantee that the conclusion is true. However, there is a lesser kind of guarantee that this premise may provide. If the argument is a strong one, the premise makes

the conclusion *probable*; it provides a *good reason* for thinking the conclusion is true; it makes the conclusion *plausible*. Instead of an absolute guarantee, we have here a weaker guarantee. You are running a risk of being wrong about your conclusion, even if your premise is true. But this risk might be a reasonable one to take. The conclusion might be a good bet, given that the premise is true.

Two Gambling Strategies

The language in the previous paragraph suggests that you can think about the difference between deductive and nondeductive arguments in terms of ideas about gambling. Consider two sorts of gamblers. The first I'll call the *extreme conservative*. This individual refuses to wager unless winning is a sure thing. The second individual I'll call the *thoughtful risk taker*. This individual at times enters into risky gambles hoping to win. Each strategy has its virtue and its limitation. The virtue of the conservative strategy is that you'll never lose a gamble. Its limitation is that there are gambles you will decline to take that you could have won. The limitation of thoughtful risk taking is that you can lose money. Its virtue is that it can lead you to win wagers by taking risks.

Limiting yourself to deductively valid arguments is a conservative strategy. You avoid the risk of reaching false conclusions from true premises. The limitation is that you decline to say anything that goes beyond the evidence. Nondeductive arguments are riskier. The gain is that you can reach true conclusions that go beyond what the premises say; the risk is that you may draw a false conclusion from true premises.

In science as well as in everyday life, we make nondeductive inferences all the time. We are often prepared to take risks. Each of us has beliefs about the future. These, however, aren't deduced from the observations we made in the present and past.

Universal Laws

Science is a risky business in another way. Scientists often try to reach conclusions about *universal laws*. An example is Isaac Newton's (1642–1727) universal law of gravitation, which you may have studied in high school. This law says that the gravitational attraction between two bodies is proportional to the products of their masses and is inversely proportional to the square of the distance between them. This law describes how much gravitational attraction there is between any two objects, no matter where those objects are located and no matter when those objects exist. Newton's law is *universal* in scope—it describes what is true at any time and place. This isn't an isolated example. In lots of sciences, there are universal statements that scientists think are well supported by evidence.

Could Newton have deduced his law from the observations he made and the experiments he conducted? *No*. His law is universal in scope. His observations, however, were conducted in a narrowly limited range of places and times. Newton didn't go backward in time to check if his law held true 3 million years ago. Nor did he send a spaceship to a distant galaxy to do the required measurements. When scientists conclude that a universal law is true or probably true, based on premises that describe the observations they have made, they aren't making deductively valid arguments.

Science is an ambitious enterprise. Science ventures beyond what is strictly observed in the here and now, just as the conclusion in a nondeductive argument ventures beyond the information strictly contained in the premises.

Detective Work

I said before that nondeductive arguments are constantly used both in science and in everyday life. Newton was my scientific example. Let me describe the calculations of Sherlock Holmes as my everyday one.

Holmes was constantly telling Watson that he figures out detective problems by "deduction." Although Holmes was a very good detective, I doubt that he solved his puzzles by strictly deductive methods. Holmes didn't observe the crimes he was later called upon to investigate. What he observed were *clues*. For example, suppose Holmes is trying to solve a murder. He wonders whether Moriarty is the murderer. The clues Holmes gathers include a gun, a cigar butt, and a fresh footprint, all found at the scene of the crime. Suppose the gun has an "M" carved in the handle, the cigar is Moriarty's favorite brand, and the footprint is the size that would be produced by Moriarty's ample foot. Can Holmes deduce from these clues that Moriarty is the murderer? No. Although the information may make that conclusion plausible or probable, it doesn't absolutely rule out the possibility that someone else did the dirty deed.

I've been emphasizing that in a strong nondeductive inference, the premises make the conclusion plausible or probable; they don't absolutely guarantee that the conclusion must be true. I now want to talk about the difference between two sorts of nondeductive inference—inductive and abductive.

Induction

Inductive inference involves taking a description of some sample and extending that description to items outside the sample. The voter survey discussed before provides an example:

> 60 percent of the county voters called are Democrats.
>
> ---
>
> Approximately 60 percent of the county voters are Democrats.

Notice that in this example the vocabulary present in the argument's conclusion is already used in the premise. Although the conclusion goes beyond what the premise asserts (which is what makes the argument nondeductive), no new concepts are introduced in the conclusion.

Two Factors Influence Inductive Strength

In the case of deduction, I said that an inference is either deductively valid or it isn't. Validity is a yes/no affair. It is like pregnancy. Inductive strength, however, isn't a yes/no matter; inductive arguments are either stronger or weaker. Inductive strength is a matter of degree.

Two factors affect how strong or weak an inductive argument is. The first is sample size. If you called one thousand individuals in your phone survey, that would make your conclusion stronger than if you had called only one hundred. The second factor is the representativeness or unbiasedness of the sample. If you called one thousand individuals drawn at random from a list of voters, that would make the resulting inference stronger than if you had called one thousand members of labor unions. The percentage of Democrats in labor unions may be higher than that found in the population as a whole. If so, you are biasing your sample by drawing it exclusively from a union membership list.

By making a telephone survey, you are failing to contact people who don't have phones. Is this a problem? That depends. If the percentage of Democrats with phones is approximately

equal to the percentage of Democrats among registered voters, no bias is introduced. On the other hand, if people with phones are disproportionately Democrats or disproportionately Republicans, your phone survey will have introduced a bias.

How do you avoid having a biased sample? Sometimes this is done by "randomization." If you have a list of all the county voters, drawing names "at random" means that each name has the same chance of being selected. However, this process of selecting at random can fail to ensure an unbiased sample. For example, suppose you draw names at random, but all the people you contact happen to be women. If women are disproportionately Democrats or disproportionately Republicans, your sample is biased. I don't say that random draws from the voter list will *probably* result in this sort of bias. My point is just that randomizing doesn't absolutely guarantee that your sample is unbiased. I won't say more here about how you can avoid having a biased sample. This fine point aside, the basic idea is this: Inductive arguments are stronger or weaker according to (1) the sample size and (2) the unbiasedness of the sample.

Abduction

I now move to abduction—inference to the best explanation. I'll begin with an example of an abductive inference that was important in the history of science. After saying what is distinctive about this form of inference, I'll describe two principles that are relevant to deciding whether an abductive inference is strong or weak.

Inferring What Isn't Observed

Gregor Mendel (1822–1884) was an experimental biologist who worked in a monastery in Moravia. He is credited with having discovered genes, the particles in living things that allow parents to transmit characteristics to offspring in reproduction.

The first thing to note about Mendel's discovery is that he never actually saw a gene (or an "element," as Mendel called them). Although more powerful microscopes made this possible later, Mendel never saw even one of them. Rather, Mendel reasoned that the observations he made could be explained if genes existed and had the characteristics he specified.

Mendel ran breeding experiments in the monastery's garden. He crossed tall pea plants with short ones and noted the proportion of tall and short plants among the offspring. Similarly, he crossed plants that had wrinkled peas with plants that had smooth peas, and noted the mix of wrinkled and smooth plants among the progeny. He then crossed some of those offspring with each other, and observed the proportion of various characteristics found in the next generation.

Mendel saw that when plants of certain sorts are crossed, their offspring exhibit characteristics in very definite proportions. Mendel asked himself a question that never figured in my discussion of induction. He asked *why* the crosses produced offspring with characteristics distributed into such proportions.

This why-question led Mendel to invent a story. He said, suppose each plant contains particles (genes) that control the observed characteristics of tall and short, wrinkled and smooth, in certain specific ways. He conjectured that each parent contributes half its genes to the offspring and that this process occurs in accordance with definite rules. The whole invented story had this property: If the story were true, that would explain why the breeding experiments had the results that Mendel observed them to have.

It should be quite clear that Mendel's theory of the gene went beyond the observations then available to him. He never saw a gene, but his theory postulates the existence of such

things. I noted before that it is a general feature of nondeductive inference, whether inductive or abductive, that the conclusion goes beyond the premises. We see here, however, a respect in which abduction differs from induction.

Abduction Differs from Induction

If Mendel had made an inductive inference, he simply would have claimed that the observed results of the experiments he ran in his garden would also occur elsewhere. His experiment was made in Europe at the end of the nineteenth century. An inductive extension of the description of his experiment might conclude that the same results would occur in twentieth-century North America as well. Had Mendel limited himself to this suggestion, no one would remember him now as the father of genetics. His important inference was abductive, not inductive. He didn't simply claim that the experiment could be replicated. Rather, he formulated a theoretical explanation of *why* the results occurred. Mendel's inference drew a conclusion concerning something he did *not* see (genes), based on premises that described what he *did* see (the results of the experimental crossings).

Can You Deduce the Explanation from the Observations?

Let's attend to Mendel's inference more carefully. The following is *not* a deductively valid argument:

> Experimental crosses in the pea plants were observed to exhibit such-and-such results.
>
> ---
>
> There are genes, and they obey laws *L*.

Remember: You can't validly deduce a theory from a set of observations.

Why can't you do this? Basically, the last argument attempts to infer a theory about the cause from the observation of its effects. There are, however, lots of possible causes that might have been responsible for the observed effects. The argument is deductively invalid for the same reason the following argument is also invalid:

> A pistol with an "M" on the handle, an El Supremo cigar butt, and a size 12 footprint were found next to the murder victim's body.
>
> ---
>
> Moriarty is the murderer.

Although Moriarty may be the most plausible suspect, the clues, in themselves, don't absolutely guarantee that he must have done the deed. Mendel and Holmes were making an inference about what is *probably* true, given the observations. They weren't inferring what is *absolutely guaranteed* to be true by the observations.

Deducing Observational Predictions from a Theory

If a set of observations doesn't deductively imply a theory, then perhaps the reverse is true: Maybe a theory deductively implies some observations. This corresponds more closely to what Mendel did. He saw that his theory of the gene implies that certain experimental results ought to occur. He then saw that those predictions came true. He concluded that the truth of the predictions was evidence that the theory is true.

When the Prediction Comes True

So a better representation of Mendel's inference might go like this. The theory entailed a prediction. The prediction came true. Hence, the theory is probably true. What we now need to see is that the following form of argument is not deductively valid:

> If there are genes and they obey laws *L*, then experimental crosses in the pea plants should exhibit such-and-such results.
> Experimental crosses in the pea plants were observed to exhibit such-and-such results.
> ---
> There are genes and they obey laws *L*.

This is deductively invalid for the same reason that the following argument is too:

> If Jones lives in Wisconsin, then Jones lives in the United States.
> Jones lives in the United States.
> ---
> Jones lives in Wisconsin.

Note that these two arguments have the same logical form.

Scientists often test their theories by seeing whether the predictions made by the theories come true. There is nothing wrong with doing this. The point, however, is that the truth of the theory doesn't follow deductively from the truth of the prediction. Scientists are reasoning *non*deductively when they decide that a theory is plausible because its predictions have come true. *Successful prediction isn't absolutely conclusive proof that the theory is true.*

When the Prediction Turns Out to Be False

On the other hand, if the predictions entailed by Mendel's theory had come out false, that would have allowed him to deduce that the theory is mistaken. That is, the following argument *is* deductively valid:

> If there are genes and they obey laws *L*, then experimental crosses in the pea plants should exhibit such-and-such results.
> Experimental crosses in the pea plants did not exhibit such-and-such results.
> ---
> It is false that there are genes and they obey laws *L*.

In other words, *a failed prediction is conclusive proof that the theory that deductively entails the prediction is false.*

How True Predictions and False Predictions Are Interpreted

Let's generalize these points. Let *T* be a theory and *P* be a prediction the theory makes. If the prediction comes out true, we cannot deduce that the theory is true. If, however, the prediction comes out false, we can deduce the theory is false:

Invalid	*Valid*
If *T*, then *P*	If *T*, then *P*
P	Not-*P*
---	---
T	Not-*T*

3 Deducing that a Theory is true

Recall from Chapter 2 that a deductively invalid argument can be turned into a valid one by adding premises. I will now exploit this fact to show how the truth of a theory can be deduced from the fact that it makes a successful prediction, *if certain further assumptions are made*.

Suppose we wish to design an experiment that tests two theories. Here the problem isn't one of evaluating a single theory, but of seeing which of two theories is more plausible. To test one theory (T_1) against another (T_2), we want to find a prediction over which they disagree. Suppose T_1 predicts that P will be true, while T_2 predicts that P will be false. If we assume that one or the other theory is true, we can find out whether P comes true and then deduce which theory is true. For example, if P turns out to be true, we can reason as follows:

T_1 or T_2
If T_1, then P.
If T_2, then not-P.
P

T_1

Notice that this argument is deductively valid. Also note that if P had turned out to be false, we could have deduced that T_2 is correct.

The difference between these two arguments—one deductively valid, the other not—suggests there is an important difference between the way scientists argue that theories are true and the way they argue that theories are false. It is possible to reject a theory just on the basis of the false predictions it makes, using a deductively valid argument; but it isn't possible to accept a theory just on the basis of the true predictions it makes, using a deductively valid argument. I will discuss this difference again in Chapter 8.

So far I've explained how a deductively valid argument can lead scientists to reject a suggested explanation. But how do scientists ever interpret observations as providing strong evidence in favor of the explanations they consider? This must involve a nondeductive inference. But what are the rules that govern such inferences? I now present two ideas that are relevant to evaluating abductive arguments. I call these the *Surprise Principle* and the *Only Game in Town Fallacy*.

The Surprise Principle: When Does Successful Prediction Provide Strong Evidence?

I've argued that you can't validly deduce that a theory is true just from the fact that some prediction it makes comes true. But maybe only a small modification of this idea is needed. Perhaps all we need to say is that a theory is made highly probable or plausible when a prediction it makes comes true. I now want to explain why this reformulation is also mistaken.

An unconscious patient is brought into the emergency room of a hospital. What is wrong? What would explain the fact that the patient is unconscious? The doctor on duty considers the hypothesis that the patient is having a heart attack. How should the doctor test whether this hypothesis is true? Well, the hypothesis predicts that the patient will have a heart (after all, if someone is having a heart attack, he or she must have a heart). The doctor verifies that this prediction is correct—the patient, indeed, does have a heart. Has the doctor