Core Questions in Philosophy

A Text with Readings

Fifth Edition

Elliott Sober

University of Wisconsin–Madison

4/'08
For Joel Velasco,
Thanks for your
help on this.
Elliott

PEARSON
Prentice Hall

Upper Saddle River, New Jersey 07458

PART I
INTRODUCTION

CHAPTER 1

What Is Philosophy?

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 they believe. 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 a fruit. Everyday 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 "What is philosophy?" 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 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. 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. Some people long ago 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 also is 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—he will have a true belief. But it isn't plausible to say that he knew, on those races about which he turned out to be right, that the first 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 are true. Knowledge, however, we never have. We don't even know those things 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 two 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 if our personalities are the results of factors outside our control.

The last problem area I'll 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 world really is. The question about ethics is whether both these realms (subjective and objective) exist in ethics, or if 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? *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 on 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 in fact, 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 senses (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; and geologists have sought to find out if the continents rest on movable plates. However, none of these sciences really bothers with the question of why we should think there are physical objects. The various sciences simply *assume* 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 the 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. So philosophers, as well as scientists, do 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 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 helps make 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, 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 degrees).

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 everyday 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?

A Problem for Further Thought

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

CHAPTER 2

Deductive Arguments

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 good 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 provide you a good reason 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 also should 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:

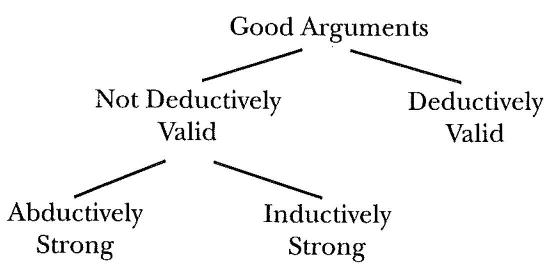

I'll treat the three categories (deductively valid, inductively strong, and abductively strong) as *mutually exclusive.* If an argument belongs to one category, it can't belong to any of the others. At the end of Chapter 3, I'll modify this classification slightly.

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. Though 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 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, though 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. Though 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* are 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 U.S.A.
Sam lives in the U.S.A.

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 California.

TESTING FOR INVALIDITY

So here's a strategy to use if you want to know whether an argument is invalid: First, ignore the content and isolate the logical form (the "skeleton") of the argument. Get rid of the distracting subject matter since that is irrelevant; what you want to focus on is the logical form. 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, since the conclusion is false. But 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 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 make it valid? There is. By adding premises, you can turn a deductively invalid argument into a valid one. Consider the following argument:

Smith lives in the U.S.A.

Smith lives in Wisconsin.

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

Smith lives in the U.S.A.
Everyone who lives in the U.S.A. lives in Wisconsin.

Smith lives in Wisconsin.

Notice that the conclusion now follows from the premises. The trouble is 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's invalidity, you have to criticize its premises for not being true. The following argument pair is different. Here one can repair the defect of invalidity and obtain a perfectly fine argument. Notice first that the following argument isn't deductively valid:

Smith lives in Wisconsin.

Smith lives in the U.S.A.

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

Smith lives in Wisconsin.
Everyone who lives in Wisconsin lives in the U.S.A.

Smith lives in the U.S.A.

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 are true as well.

This idea will come up repeatedly when I discuss various philosophical arguments. I sometimes will 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 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 so 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.

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 the conclusion is true. If an argument is deductively valid and has true premises, is that sufficient to make it good? To see why validity and true premises aren't enough, consider the following argument:

Lemons are yellow.
____
Lemons are yellow.

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. 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 U.S.A." This is equivalent in meaning to saying that "If you don't live in the U.S.A., 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 U.S.A." is true, but "If you live in the U.S.A., then you live in Wisconsin" is false. These two if/then statements can't mean the same thing, since what the one says is true while what the other says is false. "If *Q*, then *P*" is termed the *converse* of the conditional "If *P*, then *Q*." A conditional and its converse are not equivalent.

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?

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 also should see if the argument begs the question.

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

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 this is true is that the Bible says there is a God. Would this argument convince someone who didn't already believe that there is a God? Probably not. Anyone who doubts 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.

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.

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 it? 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 may 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 radical 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.

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 my own mind (for example, "I am in pain") are made true just by my believing they are true (Chapter 13).
>
> Ethical statements are true just because God, society, or some individual agent thinks they are (Chapter 31).

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 count as 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)	(b)
If *P*, then *Q*	If *P*, then *Q*
P	*Q*
———	———
Q	*P*

(c)	(d)
If *P*, then *Q*	If *P*, then *Q*
Not-*P*	Not-*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 14.
6. A sign on a store says "No shoes, no service." Does this mean that if you wear shoes, then you will be served?
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?

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, however. 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 U.S.A.
   ———
   Fred lives in the U.S.A.

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

(a)	(b)
X	*R*
If X, then Y	*S*
———	———
Y	*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.

CHAPTER 3

Inductive and Abductive Arguments

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. If the premises are true, the conclusion can't 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 also is 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 some 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, 1000 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 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% of the people called said they are Democrats.
> ________________
> 60% of the people called are Democrats.

> 60% of the people called are Democrats.
> ________________
> Approximately 60% of the voters in the county are Democrats.

Why can't you deduce these things validly? 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 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 often are 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. In saying this, though, 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 deductive 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 reach a false conclusion from true premises.

PREDICTIVE BELIEFS

In science as well as in everyday life, we make nondeductive inferences all the time. We often are 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 rather 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 a very 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 in both 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

I'll begin with some properties of 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% of the county voters called are Democrats.
> ———————
> Approximately 60% 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; 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 1,000 individuals in your phone survey, that would make the conclusion stronger than if you had called only 100. The second factor is the representativeness or unbiasedness of the sample. If you called 1,000 individuals drawn at random from a list of voters, that would make the resulting inference stronger than if you had called 1,000 members of labor unions. The percentage of Democrats in labor unions may be higher than that 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.

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.

INFERRING WHAT ISN'T OBSERVED

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 he noted the mix of wrinkled and smooth plants among the progeny. He then crossed some of those offspring with each other, and he saw the proportion of various characteristics found in the next generation.

Mendel observed 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 anywhere the experiment was carried out. 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 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.

Though 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 U.S.A.
> Jones lives in the U.S.A.
> ______________
> 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 didn't 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 implying the prediction is false.*

HOW TRUE PREDICTIONS AND FALSE PREDICTIONS ARE INTERPRETED

Let's generalize these points. Let *T* be a theory and *P* a prediction the theory makes. If the prediction comes out true, we can't 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*

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 now will exploit this fact to show how the truth of a theory can be deduced from the fact that it makes a successful prediction, *if you make certain further assumptions.*

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. discuss this difference again in Chapter 8.

So far I've explained how a deductively valid argument can lead a scientist 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. These I call 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 also is 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 thereby obtained strong evidence that the patient is having a heart attack? Clearly not. This is an example in which you don't obtain serious support for a hypothesis just by showing that one of its predictions is correct.

You go to the gym, and someone tells you he is an Olympic weight lifter. You reason that if he is an Olympic weight lifter, he should be able to pick up the hat you are wearing. You offer him your hat; he lifts it without difficulty. Have you thereby obtained strong evidence that he is an Olympic weight lifter? Clearly not. Once again, the hypothesis under test isn't strongly supported by the fact that one of its predictions turns out to be correct.

What has gone wrong in these two cases? In the first example, the presence of a heart isn't strong evidence that the patient is having a heart attack. The reason is that you would expect the individual to have a heart even if he weren't having a heart attack. In the second example, the man's lifting the hat isn't strong evidence that he is an Olympic weight lifter. The reason is that you would expect him to be able to do this even if he weren't an Olympic lifter.

What should you look for if you want to test the hypothesis that the patient has had a heart attack? What you want to find is a symptom that you would expect to find if the patient were having a heart attack, *but would expect to not be present if the patient were not having a heart attack.* Suppose an erratic EKG (electrocardiogram) almost always occurs when there is a heart attack but rarely occurs when there is no heart attack. This means that the presence of an erratic EKG would be strong evidence that the patient is suffering a heart attack.

What sort of test should you use if you want to see if the man in the gym really is an Olympic weight lifter? Suppose Olympic weight lifters (of this man's weight) can almost always lift 400 pounds, but people who aren't Olympic lifters can rarely do this. This means that his managing to lift 400 pounds would be strong evidence that he is an Olympic weight lifter.

Think of the unconscious patient and the weight lifter as posing *discrimination problems.* The problem is to find evidence that strongly discriminates between two hypotheses. In the first example, the competing hypotheses are:

H_1: The patient is having a heart attack.
H_2: The patient isn't having a heart attack.

An erratic EKG strongly favors H_1 over H_2; the mere fact that the patient has a heart does not.

The same holds true for the second example. The problem is to find evidence that discriminates between the following two hypotheses:

H_1: This man is an Olympic weight lifter.
H_2: This man isn't an Olympic weight lifter.

The fact that the man can lift 400 pounds strongly favors H_1 over H_2; the fact that he can lift a hat does not.

The Surprise Principle describes what it takes for an observation to strongly favor one hypothesis over another:

> *The Surprise Principle:* An observation O strongly supports H_1 over H_2 if both the following conditions are satisfied, but not otherwise: (1) if H_1 were true, O is to be expected; and (2) if H_2 were true, O would have been unexpected.

Let's apply this principle to the first example. Consider the observation that the patient has an erratic EKG. If the patient were having a heart attack (H_1), we would expect him to have an erratic EKG. And if the patient weren't having a heart attack (H_2), we would expect him not to have an erratic EKG. This explains why the erratic EKG is strong evidence of a heart attack—the EKG strongly favors H_1 over H_2. Now consider the observation that the patient has a heart. If the patient were having a heart attack, we would expect him to have a heart. But we would expect him to have a heart even if he weren't having a heart attack. This explains why the presence of a heart isn't strong evidence that the patient is having a heart attack—the observation

No Surprise/Surprise

The Surprise Principle involves two requirements. It would be more descriptive, though more verbose, to call the idea the No Surprise/Surprise Principle.

The Surprise Principle describes when an observation O strongly favors one hypothesis (H_1) over another (H_2). There are two requirements:

(1) If H_1 were true, you would expect O to be true.

(2) If H_2 were true, you would expect O to be false.

That is, (1) if H_1 were true, O would be unsurprising; (2) if H_2 were true, O would be surprising.

The question to focus on is *not* whether the hypotheses (H_1 or H_2) would be surprising. The Surprise Principle has nothing to do with this. To apply the Surprise Principle, you must get clearly in mind what the hypotheses are and what the observation is.

doesn't strongly favor H_1 over H_2. We were looking for an explanation of why the erratic EKG provides telling evidence whereas the presence of a heart doesn't. The key is to be found in condition (2) of the Surprise Principle.

Take a few minutes to apply the Surprise Principle to the example of the weight lifter. Make sure you see how the principle explains why being able to lift a hat isn't strong evidence, whereas being able to lift 400 pounds is.

EVIDENCE MAY DISCRIMINATE BETWEEN SOME HYPOTHESES WHILE FAILING TO DISCRIMINATE BETWEEN OTHERS

In the examples of the heart attack and the weight lifter, H_2 says only that H_1 is false. In other abductions, however, H_2 may say more than this. Suppose you see someone crossing campus carrying several philosophy books. You wonder whether the person is a philosophy major. Two hypotheses to consider are

H_1: The person is a philosophy major.
H_2: The person is an engineering major.

According to the Surprise Principle, the observation you have made favors H_1 over H_2. But now consider the following third hypothesis:

H_3: The person isn't a student, but is in the business of buying and selling philosophy books.

Although your observation discriminates between H_1 and H_2, it does not discriminate between H_1 and H_3.

This brings out an important fact about how the Surprise Principle applies to abductive inferences. *If you want to know whether an observation strongly supports a hypothesis, ask yourself what the alternative hypotheses are.* For an observation to strongly support a hypothesis is for it to strongly favor that hypothesis over the others with which it competes.

TRUE PREDICTION ISN'T ENOUGH

The point of the examples about the unconscious patient and the weight lifter was to show why successful prediction doesn't automatically provide strong evidence. If someone is having a heart attack, that predicts that he will have a heart; but the presence of a heart isn't strong evidence the person is having a heart attack. If someone is an Olympic weight lifter, that predicts that he will be able to lift a hat; but lifting a hat isn't strong evidence the person is an Olympic weight lifter.

There is a scene in the Monty Python movie *The Life of Brian* that illustrates this idea. The setting is a marketplace. Around the perimeter of the market are assorted prophets and soothsayers. The camera pans from one to the other. We see that in each case, the prophet holds a crowd of people in rapt attention. The first prophet predicts that tomorrow a purple monster will rise out of the desert and devour three villages. The crowd is amazed at these predictions and no doubt will conclude that the prophet has special powers to foresee the future if the predictions come true. After showing us a few prophets of this sort, the camera comes to an individual who very calmly makes the following predictions: Tomorrow, many people will get up early. Others will sleep longer. Some people will decide to have breakfast, while others will postpone eating until later in the day. And so on. The joke is that the crowd in front of this guy is just as awestruck as the crowds were in front of the more outlandish prophets.

There is a lesson here. How would you test the hypothesis that someone has special powers to foresee the future? If he predicts events that people without special powers can easily predict, it isn't very impressive that his predictions come true. If, however, he predicts events that normal people aren't able to foresee, and then these predictions come true, we are more impressed. The Surprise Principle explains why the success of "safe" predictions provides less compelling evidence than the success of "daring" predictions.

Here's a related example of the Surprise Principle in action. Many people thought that the astrologer Jeanne Dixon had special powers to predict the future. After all, she predicted the assassination of President John F. Kennedy and several other events that no one could have guessed were going to happen. If Jeanne Dixon played it safe and only predicted events that everybody knew were going to happen, we wouldn't be impressed. But aren't we being sensible in reasoning that she probably does have special powers, since this and other daring predictions have come true?

Although this reasoning may seem to conform to the Surprise Principle, it doesn't. The thing people sometimes forget is that Jeanne Dixon made thousands of predictions and most of them turned out false. It isn't surprising at all that some handful of these should have come out true. Although Jeanne Dixon predicted the Kennedy assassination, which surprised (practically) everyone when it happened, it isn't at all surprising that someone with no special powers should be lucky a few times every thousand tries. The Surprise Principle, properly understood, tells us why we shouldn't take Jeanne Dixon's few successes as strong evidence that she had special powers.

MODEST FAVORING

The Surprise Principle shows when an observation *strongly* favors one hypothesis over another. However, sometimes our observations are not so telling and unequivocal. Sometimes the observations (O) favor one hypothesis (H_1) over another (H_2), but only modestly. This will be true when H_1 confers on O a higher probability than H_2 does, but the difference is modest. What the Surprise Principle and this idea about O's modestly favoring H_1 over H_2 have in common is this: Both make use of the idea that an observation favors one hypothesis (H_1) over another (H_2). This will be true when the probability of O, according to H_1, exceeds the probability of O, according to H_2. We will examine the concept of probability in more detail in Chapter 18.

THE SURPRISE PRINCIPLE SUMMARIZED

In summary, the Surprise Principle gives advice on what a hypothesis must do if it is to be strongly supported by the predictions it makes. First, the hypothesis shouldn't make false predictions. Second, among the true predictions the hypothesis makes, there should be predictions we would expect not to come true if the hypothesis were false.

When we ask whether an observation O strongly supports some hypothesis H_1, the Surprise Principle requires that we specify what the alternative hypotheses are against which H_1 competes. It may turn out that O strongly favors H_1 over H_2, but that O doesn't strongly favor H_1 over H_3. The example of the person carrying philosophy books across campus illustrates this point.

THE ONLY GAME IN TOWN FALLACY

I turn now to a second principle for evaluating abductive inferences. Suppose you and I are sitting in a cabin in the woods. We hear a strange rumbling sound in the attic. You ask, "I wonder what that could be?" I reply, "That is the sound of gremlins bowling in the attic." You, being a sensible person, reply, "I really don't think there are gremlins in the attic." I then challenge you to produce a more plausible explanation of the noises. You reply "Gosh, I really don't have any idea why those noises occurred.

An Investment Swindle

Suppose you received a letter every month for a year from an investment firm. In each letter, a prediction is made as to whether "the stock of the month" will increase or decline in value during the next 30 days. You keep track of what happens to the stocks described each month. Each prediction comes true. Would you conclude from this that the investment firm has a method for reliably predicting stock market events?

Some years ago an "investment firm" sent out such letters, but the mailing was a swindle. The firm began with a list of 10,000 investors. In the first month, 5,000 investors received letters predicting that stock *A* would go up; the other 5,000 received letters saying that stock *A* would go down. The firm then waited to see which prediction came true. During the second month, the firm sent letters to the 5,000 people who had received a true prediction during the first month. In the second month, 2,500 investors received letters predicting that stock *B* would go up; the other 2,500 received letters saying that *B* would go down. The process was repeated, so that by the end of 10 months a small number of investors had received 10 letters, each containing a successful prediction.

The company then wrote to those people, asking each to invest a large sum of money. Most did so. The company then absconded with the funds. (This story is from Daniel Dennett's *Brainstorms: Philosophical Essays on Mind and Psychology*, Cambridge, Mass., MIT Press, 1978.)

The investors who were swindled thought they were making a reasonable abductive inference on the basis of the company's track record. Describe the premise, the conclusion, and the reasoning that led the investors to think the conclusion was highly plausible. Were the investors making a strong inference?

I just think your story is implausible." To this humble admission on your part, I make the following rejoinder: "Look, my story, if true, would explain why we just heard those strange noises. If you don't want to accept my explanation, you must produce a more plausible explanation of your own. If you can't, you have to accept my explanation of the noises."

What I just did was commit an abductive fallacy (mistake in reasoning), which I'll call the *Only Game in Town Fallacy*. The fact that you can't think of a more plausible explanation of the noises doesn't oblige you to accept the story I constructed. There is an alternative, which is simply to admit that the noises are something you don't know how to explain.

Abduction is sometimes described loosely as follows: If a theory explains some observation, and if no rival account is available that can do a better job of explaining it, then you should accept the theory. Although this description of abduction is roughly correct, it makes the mistake of sanctioning the Only Game in Town Fallacy. The fact that no rival account is better than the explanation I construct doesn't show my explanation is even minimally plausible. My gremlin theory is pretty silly, although maybe there is nothing now available that is clearly superior to it.

I won't at this point try to fine-tune the idea of abductive inference any further, even though there is lots more to be said about it. We now have before us the basic idea of inference to the best explanation. We've seen that it is an important part of the scientific method. And I've described two principles that help guide us in evaluating whether an abductive inference is strong or weak.

Review Questions

1. What is the difference between deductive validity and inductive strength?
2. What is the difference between induction and abduction?
3. What factors affect how strong an inductive argument is?
4. Suppose a given observation discriminates between two hypotheses, but a second observation fails to do this. Construct an example, different from the ones presented in this chapter, illustrating the point. Show how the Surprise Principle applies to your example.
5. An observation can succeed in discriminating between hypotheses H_1 and H_2 but fail to discriminate between H_1 and H_3. Construct an example that illustrates this point that is different from the ones presented in this chapter. Show how the Surprise Principle applies to your example.
6. What is the Only Game in Town Fallacy? What does it mean to call it a "fallacy"?

Problems for Further Thought

1. Suppose you wanted to find out what percentage of the adults in your county are vegetarians. You obtain a list of unmarried adults in the county and contact them to do a survey. Is this sample a biased one? Why or why not?
2. Although induction and abduction were described in this chapter as separate kinds of inference, they have a good deal in common. The Surprise Principle was introduced as applying to abduction, but it applies to induction as well. Suppose an urn is filled with 1,000 balls, each of them either red or green. You reach into the urn, sampling at random, and bring out 100 balls. Your sample contains 50 red balls and 50 green ones. Here are some hypotheses to consider:

   H_1: All the balls in the urn are green.
   H_2: 75% of the balls are green.
   H_3: 50% of the balls are green.
   H_4: 25% of the balls are green.

   Suppose you think that H_1 and H_2 are the only possibilities. Does the observation strongly support one over the other? Which? How does the Surprise Principle apply to this question? Suppose, instead, that you think that H_2 and H_3 are the only possibilities. Does the sample strongly support one over the other? If so, which? Suppose, finally, that you think that H_2 and H_4 are the only possibilites. Which is best supported? Why?

3. Why think that any of the beliefs you have about the world outside your own mind are true? For example, why are you now entitled to think there is a printed page in front of you? Presumably you believe this on the basis of sense experience (sight, touch, etc.). Construct an abductive argument whose conclusion is that there is a printed page in front of you. Make sure your inference obeys the Surprise Principle. Does this abductive argument prove there is a printed page in front of you? Explain.
4. At the beginning of Chapter 2, I presented deductive validity, inductive strength, and abductive strength as mutually exclusive categories. This means that if an argument belongs to one category, it doesn't belong to any of the others. This is generally correct, but not always. Here is an inductive argument that isn't deductively valid:

   I've observed 1,000 emeralds and all have been green.
   ―――――――――
   All emeralds are green.

   However, by adding a premise, I can produce a deductively valid argument:

   I've observed 1,000 emeralds and all have been green.
   If there are over 500 emeralds in the universe, they will all have the same color.
   ―――――――――
   All emeralds are green.

   Both of these arguments are inductive in the sense that both involve drawing a sample from a population and reaching a conclusion about that population. If so, some inductively strong arguments are also deductively valid.

   The same point holds for abductively strong arguments. Usually they aren't deductively valid, but sometimes they are. Construct a strong abductive argument that obeys the Surprise Principle. Show how it can be made deductively valid by adding or modifying a premise.
5. There is a difference between *not expecting O* and *expecting not-O*. A person who never considers whether *O* is true does not expect *O*, but it would be wrong to say that she expects not-*O*. Of course, if someone expects not-*O*, it will also be true that she does not expect *O*. Thus, "*S* expects not-*O*" implies "*S* does not expect *O*," but the reverse is not true. With this logical point in mind, explain why condition 2 of the Surprise Principle is formulated by saying "H_2 leads you to expect not-*O*," rather than saying that "H_2 does not lead you to expect *O*."

Suggestions for Further Reading

ON THE NATURE OF PHILOSOPHY

The essays in Charles Bontempo and S. Jack Odell (eds.), *The Owl of Minerva: Philosophers on Philosophy,* New York, McGraw-Hill, 1975.

Bertrand Russell, "The Value of Philosophy," in *The Problems of Philosophy,* Oxford, Oxford University Press, 1950.

ON DIFFERENT FORMS OF ARGUMENT

Monroe Beardsley, *Thinking Straight,* Englewood Cliffs, New Jersey, Prentice Hall, 1975.

Irving Copi, *Informal Logic,* New York, Macmillan, 1986.

Robert Fogelin, *Understanding Arguments,* New York, Harcourt, Brace, Jovanovich, 1978.

Ronald Giere, *Understanding Scientific Reasoning,* New York, Holt, Rinehart, and Winston, 1984.

Merrilee Salmon, *Logic and Critical Thinking,* San Diego, Harcourt, Brace, Jovanovich, 1984.

PART II
PHILOSOPHY OF RELIGION

CHAPTER 4

Aquinas's First Four Ways

Saint Thomas Aquinas (c. 1224–1274) was an enormously accomplished theologian and philosopher. In his masterwork, the *Summa Theologiae,* he presents five proofs that God exists (Aquinas called them "the five ways"). I'll discuss the first four now and the fifth in the next chapter.

Each of Aquinas's arguments begins with a simple observation that is supposed to be obvious to everyone. For example, the first argument begins with the observation that physical objects are in motion. Each argument then proceeds through various other premises to the conclusion that the explanation of the initial observation is that there is a God. Aquinas intends each of his proofs to be deductively valid.

In Chapter 3, I stressed that most of the hypotheses that scientists are interested in testing can't be deduced from observations. For example, Mendel couldn't deduce the existence and characteristics of genes from the observations he made on his pea plants. The same is true in Aquinas's arguments, as he fully realizes: You can't deduce the existence and characteristics of God just from the simple observations with which his arguments begin. The existence of motion doesn't, all by itself, deductively imply the existence of God. Aquinas's arguments always include additional principles. It is these further principles that are supposed to link the starting observation with the conclusion that God exists.

THE CONCEPT OF GOD

Before describing Aquinas's arguments, I need to say something about what he means by "God" and how I'll use that term. Aquinas took God to be a person—one who is all-powerful (omnipotent), all-knowing (omniscient), and entirely good (omnibenevolent)—all-PKG, for short. This conception of God is a familiar one in the traditions of Judaism, Christianity, and Islam (though there is room to debate, in these traditions, whether God is properly described in this way and other religions have other views about what characteristics God has). I'll assume provisionally that God, if such a being exists, has the three characteristics just mentioned. If we don't start with some preliminary picture of what God is, we won't know what we are talking about when we ask whether God exists. However, it is important to bear in mind that this conception of God is not the only one that is possible. Indeed, in Chapter 11, I'll consider an argument that suggests that God, if there is such a being, can't be all-PKG. The definition of God as an all-PKG being is a useful place to begin discussion, but it is only a point of departure.

Another caveat I should mention is that my versions of Aquinas's arguments won't be accurate in all respects. It is often a subtle historical question what this or that philosopher had in mind in a given text. In this case and later in this book, when I discuss the ideas of other philosophers, I often will examine somewhat simplified versions of the arguments they constructed. You may well ask: Why is it worthwhile studying simplified versions of a great philosopher's arguments? Admittedly, there is a loss, but there also is a gain. The main point to be made here, at the beginning of an introductory text, is that it is useful to evaluate these simpler arguments before more subtle arguments are addressed. At any rate, there is ample philosophical material to think about in the arguments I'll describe, even if these arguments don't capture the thoughts of various great thinkers with total accuracy and completeness.

THE FIRST TWO ARGUMENTS: MOTION AND CAUSALITY

Aquinas's first argument for the existence of God is the argument from motion. Here it is, formulated as a deductive argument:

(I) (1) In the natural world, there are objects that are in motion.
(2) In the natural world, objects that are in motion are always caused to move by objects other than themselves.
(3) In the natural world, causes must precede their effects.
(4) In the natural world, there are no infinite cause/effect chains.

---

(5) Hence there is an entity outside of the natural world (a *supernatural* being), which causes the motion of the first moving object that exists in the natural world.

---

(6) Hence, God exists.

Aquinas's second argument generalizes the ideas found in the first. Whereas the first argument is about motion in particular, the second argument is about causality in general:

(II) (1) The natural world includes events that occur.
(2) In the natural world, every event has a cause, and no event causes itself.
(3) In the natural world, causes must precede their effects.
(4) In the natural world, there are no infinite cause/effect chains.

---

(5) Hence there is an entity outside of nature (a supernatural being), which causes the first event that occurs in the natural world.

---

(6) Hence, God exists.

In both these arguments, I've drawn two horizontal lines to indicate that (5) is supposed to follow from premisses (1)–(4) and that (6) is supposed to follow from (5).

What do premises (2) and (3) mean in these two arguments? Let's begin with an object (O_1) that is in motion now. Since it is in motion, there must be an earlier object—call it O_2—that sent O_1 into motion. If O_2 was itself in motion, O_2 must trace back to a previous mover, O_3, and so on. In the second argument, the subject is causality, not motion, but the idea is basically the same.

Premise (4) says that there can't be a cause/effect chain that extends infinitely far back into the past. The idea is that cause/effect chains (or mover/movee chains) leading from the present back into the past have a finite number of links. The arrow in the following diagram represents the relationship of causality:

Past   Present   Future

$$O_f \longrightarrow \cdots \longrightarrow O_3 \longrightarrow O_2 \longrightarrow O_1 \longrightarrow \quad \cdots$$

Although Aquinas's argument focuses on chains that extend from present back into the past, his principle (4) also has implications about chains that extend from the present forward into the future. These also must have a finite number of links. Cause/effect chains must be finite in both directions.

I'll begin with a criticism aimed just at argument I; after that, I'll lump the two arguments together and formulate some objections that apply to both.

AQUINAS ON THE CAUSE OF MOTION

Aquinas thinks (premises 1–2) that if an object is in motion, it must be caused to move by something outside itself. Aquinas got this idea from Aristotle's physics. Aristotle (384–322 B.C.) held that if an object continues to move, its motion must be sustained by a force that keeps it in motion. If you remove the force, the object stops moving.

This idea didn't survive into modern physics. You may remember from your high school physics course that Newton, in the seventeenth century, held that an object remains in constant uniform motion unless acted on by a force. Recall that one of Newton's laws of motion is $F = ma$. This means that if an object of mass m is acted on by a force of value *F*, then it will accelerate to degree *a*. This Newtonian law says that an object that isn't acted on by a force won't accelerate, which means it will remain at rest *or in uniform motion*. Newton's laws say that an object can remain in motion forever without there being any force that sustains its motion.

Newton's laws do not exclude the following possibility: The universe contains exactly one physical object, which always moves in uniform motion without any forces ever acting on it. Of course, Newton's laws don't say that the universe we live in is like this. Our universe obviously contains more than one material object. My point, though, is that Newton's theory of motion was different from Aristotle's and so was different from Aquinas's. Aristotle and Aquinas thought that motion requires an outside force; Newton and more modern physical theories hold that it is acceleration, not simply change in position, that requires a force.

It is not difficult to rescue Aquinas's first argument from this Newtonian objection. Just replace his talk of motion with the concept of acceleration. If objects accelerate there must be a force that causes them to do so. Then we are led by the same line of reasoning to the conclusion that there must exist a supernatural entity that causes the first accelerating object in nature to accelerate.

GOD IS A PERSON, NOT JUST A CAUSE THAT EXISTS OUTSIDE OF NATURE

I turn now to some problems that the two arguments share. First, it is important to see that proposition (S)—that there is an entity outside of nature that causes the first moving object in nature to move, or that causes the first event in nature to occur—does not guarantee the existence of God, where God is understood to mean a person with something like the three properties of omnipotence, omniscience, and omnibenevolence (an all-PKG being). As Aquinas himself realized, conclusion (6) does not follow from proposition (5), in either argument.

THE BIRTHDAY FALLACY

Another problem arises when we ask whether the argument shows that there is *precisely one* first cause, or instead shows only that there is *at least one*. Suppose we grant that each causal chain in nature has a first member. According to Aquinas, each of these first members must be caused by some event outside of nature. However, it does not follow that there is exactly one such event outside of nature that set all causal chains in the natural world in motion. Here it is important to see the difference between the following two propositions; the first is different from and does not deductively imply the second:

> Every event in the natural world traces back to an event that occurs outside nature.
>
> There is a single event outside of the natural world to which each event in nature traces back.

The difference here parallels the logical difference between the following two propositions, the first of which is true and the second false:

> Every person has a birthday—a day on which he or she was born.
> There is a single day that is everybody's birthday.

I want to give a name to the mistaken idea that the second proposition follows from the first. I'll call it the *Birthday Fallacy*.

So one problem with arguments I and II is that they don't show there is exactly one first cause or unmoved mover; at best, they show that there is at least one. To think otherwise is to commit the Birthday Fallacy.

WHY CAN'T NATURE BE INFINITELY OLD?

Another objection to Aquinas's first two arguments is his claim that cause/effect chains cannot extend infinitely far into the past. Why is this impossible? If the natural world were infinitely old, each event could be caused by an earlier event. Every event that occurs in nature could have a cause that also existed in nature, so there would be no reason to infer that something outside of nature must exist as the cause of what occurs inside.

Aquinas thinks he has an answer to this question. He doesn't simply *assume* that causal chains extending into the past must be finite in length; he has an *argument* that he thinks shows why this is so. Here is his argument, which I'll reconstruct in terms of an example of a present event—your reading this page now:

> You are reading this page now.
>
> A causal chain that extends from this present event infinitely into the past, by definition, lacks a first member.
>
> If a causal chain lacked a first member, then all subsequent events in the chain could not occur.
>
> ---
>
> Hence, the causal chain leading up to your reading this page now must be finitely old.

The third premise is where this argument goes wrong. Even if we assume that no event in nature can happen without its having a cause, it does not follow that there has to be a first natural event.

Many traditional theists will agree that the world could have an infinite future—that it could go on forever. However, if an infinite future is possible, why is an infinite past ruled out? If there doesn't have to be a last event in the history of the natural world, why must there be a first? Why accept Aquinas's claim that causal chains can't extend infinitely into the past? Here Aquinas gets some help from modern physics, which views the universe as finitely old. Although it is conceivable that the universe is infinitely old, apparently there are scientific reasons to think that this isn't so.

WHY MUST EVERY EVENT IN NATURE HAVE A CAUSE?

This brings me to my last objection to Aquinas's first two arguments for the existence of God. Even if the natural world is only finitely old, why must there be an explanation of the first event that occurs in nature? That is, why must *every* event that occurs in nature have a cause?

Do scientists assume that every event has a cause? Well, it is true that scientists often try to discover the causes of events that they observe. However, this activity of searching does not require that one actually *believe* that every event has a cause. Perhaps there are exceptions to this generalization. One tries to find causal explanations for the events one observes for this reason: If you don't look, you won't discover the cause if there is one. Better to look and fail to find than never to look at all.

THE THIRD ARGUMENT: CONTINGENCY

Aquinas's third argument for the existence of God, like the first two, begins with an observation that everybody would agree is true. The observation is put in language that may be unfamiliar, but once explained, it seems clear enough. The observation is that contingent things exist. What does it mean to say that a thing is contingent? The opposite of contingency is necessity. What makes an object contingent or necessary?

NECESSARY AND CONTINGENT BEINGS

You, I, the Washington Monument, the human race, the Earth, the solar system, and the Milky Way are all examples of contingent things. Although all of us exist, we needn't have existed. The world could have failed to include any of us. If the particular sperm and egg that produced you hadn't encountered each other, you wouldn't have come into existence. Your parents would have had other children or no children at all. Likewise for the human race; although the world obviously includes human beings, it needn't have. Contingent things depend for their existence on something or other happening. In contrast, a necessary being is something that must exist no matter what. It doesn't depend for its existence on anything.

POSSIBLE WORLDS

I want to represent the concepts of necessity and contingency by introducing a new idea. Consider the totality of things that have existed, now exist, or will exist at any place in the universe. This totality comprises a giant object that I'll call "the actual

world." Imagine a census of the objects in the actual world. We might display this census on a very long time line, in which the durations of various objects are represented:

/—*w*—/

/—*x*—/ /—*y*—/ /—*z*—/

Past Now Future

Notice that the life spans, or durations, of objects *w* and *x* overlap. Objects *w* and *x* existed in the past and now are no more. Object *y* came into existence in the past, exists now, and will continue to exist until some future date, when it will cease to be. Object *z* hasn't yet come into being, but it will. Of course, this is a dreadfully incomplete census of everything that was, is, or will be. Never mind. You get the idea: The actual world consists of everything that has existed, exists now, or will exist, anywhere in the universe.

We know the world needn't have had precisely the census it does. Some of the objects in this census might have failed to exist. Other objects, which don't actually exist, might have done so. Let's call each possible way the world might have been a "possible world." There are many of these.

We now may say that an object is contingent if it exists in some but not all possible worlds. And an object is necessary if it exists in all possible worlds. Suppose the following are the censuses found in several possible worlds. As before, the horizontal line represents time within a possible world:

wx	*y*	*z*	Actual world
ab	*y*	*z*	Possible world 1
ac	*d*	*z*	Possible world 2
Past	Present	Future	

Notice that *w* exists in the actual world, but not in the two alternative possible worlds represented. On the other hand, *y* exists in the actual world and in the first possible world, but not in the second. Object *z* is different—it exists in all of the possible worlds depicted. Of course, there are more than three possible worlds, since there are more than two alternative ways the actual world might have been (the actual world, of course, is one of the possible worlds). Would *z* continue to appear if we listed not just a few possible worlds, but all of them? That is the question of whether *z* is a necessary being.

I cited some examples of contingent things—you, the Washington Monument, the Earth, and a few others. These examples might strike you as typical of everything that exists. That is, you might find plausible the following philosophical conjecture: Everything that exists is contingent. Familiar objects clearly have this property. But is *everything* contingent? Aquinas argues in his third proof of the existence of God that not everything is contingent. There is at least one necessary being—namely, God.

Here is how his proof goes:

III (1) Contingent things exist.
(2) Each contingent thing has a time at which it fails to exist ("contingent things are not omnipresent").

---

(3) So, if everything were contingent, there would be a time at which nothing exists (call this an "empty time").
(4) That empty time would have been in the past.
(5) If the world were empty at one time, it would be empty forever after (a conservation principle).

---

(6) So, if everything were contingent, nothing would exist now.
(7) But clearly, the world is not empty (see premise 1).

---

(8) So there exists a being who is not contingent.

Hence, God exists.

Aquinas's argument has two stages. First, there is his defense of proposition (6), which is contained in premises (1)–(5). Second, there is his use of proposition (6) to establish the existence of God. Before considering why Aquinas thinks (6) is correct, I want to focus on his use of (6) and (7) to infer proposition (8).

REDUCTIO AD ABSURDUM

Aquinas's proof of (8) from (6) and (7) has a distinctive logical form. His proof is a *reductio ad absurdum* argument (a "reductio," for short).

When I talked about deductive validity in Chapter 2, I emphasized that the word "valid" in logic and philosophy doesn't mean what it means in everyday life. The same holds for the idea of reducing something to absurdity. In ordinary speech this means something like "making a mockery of an idea." In logic, however, reductio arguments are perfectly good arguments—they are deductively valid.

Here's how a reductio argument works. You want to prove a proposition *P*. To do this, you argue that if *P* were false, some proposition *A* would have to be true. But you construct the argument so that *A* is an obvious falsehood (an "absurdity"). From this, you may validly conclude that *P* is true. Reductio arguments, in other words, have the following valid logical form:

If *P* is false, then *A* is true.
A is false.

P is true.

Aquinas tries to establish the existence of a necessary being by reductio. He argues that if only contingent beings existed, the world would now be empty. But it is obviously false that the world is now empty. Hence, not everything is contingent—there must exist at least one necessary being.

Let's now examine Aquinas's defense of proposition (6). Why should we think the world would now be empty if there were no necessary beings? Aquinas's reasons for thinking this are contained in premises (1)–(5).

CONTINGENCY AND ETERNITY

I'll begin by registering an objection to premise (2). Contrary to Aquinas, I submit that a contingent thing can be eternal; the fact that an object is contingent doesn't mean that there must be a time at which it fails to exist. I'll grant that familiar contingent objects aren't eternal. You are an example. You might have failed to exist; in addition, there was a time before you were born and there will be a time after you die. But this is just an example. We want to know whether *all* contingent beings must fail to be eternal.

Notice first that the considerations affecting contingency differ from those affecting the question of eternity. How can you tell whether an object is contingent in a diagram like the one given before showing the actual world and two possible worlds? You look *down* the list of possible worlds, checking to see if the object in question is present in each one. In that diagram, we might say that contingency is represented "vertically." Eternality is different. To see whether an object present in a given possible world is eternal, you look *across* the representation of that possible world, checking to see if it exists at all times in that world. So eternality is represented "horizontally."

Could an object exist in only some possible worlds and still exist at all times in the actual world? That is, could an object be both contingent and eternal? Aquinas says no. Here is a consideration that suggests he may be wrong.

The idea that material objects are made of very small indivisible particles has had a long history, going back at least to the ancient Greeks. Of course, we no longer think of atoms as indivisible particles—we have long since learned to talk about subatomic particles. Suppose for the moment, however, that the word "atom" names such indivisible particles. These particles are the basic building blocks of all material things.

One idea that has been put forward in atomistic theories is that atoms (fundamental particles) can't be created or destroyed. Large objects made of atoms can be created and destroyed by assembling and disassembling collections of atoms. But the basic particles themselves can't be destroyed, because you can't break them into pieces. Again, I'm not saying that this is true of the objects we now regard as the smallest material particles. I mention it only to describe a possible view of atoms.

Atoms, on the view I'm describing, are eternal. If a given atom exists now, it has always existed and it always will exist. Let's call one of the atoms that populates the actual world by the name "Charlie." Charlie is eternal.

Does it follow that Charlie is a necessary being—that the world couldn't have failed to include him? I would say not. The world could have contained more atoms than it does, or fewer. Indeed, it doesn't seem impossible that the world might have been entirely empty of matter. So Charlie, like all the atoms that happen to actually exist, is a contingent being.

If this is right, then Aquinas's premise (2) is mistaken. Charlie, I've claimed, is a contingent being, but he is eternal. So contingent things needn't have a time at which they fail to exist.

The argument I've given against premise (2) depends on a particular theory about the nature of atoms. If atoms really are indestructible, I can claim to have refuted Aquinas's premise. But are they, in fact, indestructible? Again, remember I'm discussing the smallest units of matter here. This is how I'm using the word "atom."

CONSERVATION LAWS IN PHYSICS

Current physics doesn't support the view that atoms are indestructible. Before Einstein, physics upheld a principle that said that the quantity of mass in the universe is constant. Einstein replaced this classical conservation law with another one. Mass isn't conserved; rather, a quantity that Einstein called "mass-energy" is conserved. This principle says that mass can be destroyed; it may be converted into energy. Modern physics, therefore, doesn't seem to allow me to tell my story about Charlie, the indestructible atom. So matter *can* be destroyed, according to current theory. But this doesn't mean that every contingently existing particle WILL be destroyed, sooner or later. It is the latter claim that Aquinas advances in premise (2). I don't know of any good reason to think this is so, although I'll leave the question open.

My conclusion thus far is that premise (2) is mistaken; a contingent thing doesn't have to have a time at which it fails to exist. Let's move on.

THE BIRTHDAY FALLACY

The transition from premise (2) to statement (3) involves a fallacy. Even if we grant that every contingent entity has a time at which it fails to exist, we can't conclude that there is an empty time. The fallacy here is the same one discussed concerning Aquinas's first two arguments. Recall the point about birthdays; "Everybody has a birthday" does not deductively imply that there is a day that is everyone's birthday.

To make this point graphic, consider the following possible world, in which no object is eternal, and yet there is no empty time:

```
←— /—1—/     /—3—/     /—5—/          —→
         /—2—/     /—4—/     /—6—/
_________________________________________
Past             Present            Future
```

The arrows at the beginning and end indicate the pattern should be repeated indefinitely into the past and into the future. So the supposition that everything is contingent doesn't imply an empty time. But there is an additional unjustified step. Even if there does have to be an empty time, why must it have been in the past? I see no reason to assume this.

NECESSARY BEINGS OTHER THAN GOD

I've noted some problems in Aquinas's attempted reductio proof that not everything is contingent. Suppose, however, that Aquinas could show that there is at least one necessary being. Could he conclude from this that God exists? I want to argue that this, too, doesn't follow. In discussing Aquinas's first two arguments, I claimed that the existence of God doesn't follow from the existence of a first cause. I now want to suggest that the existence of God doesn't follow from the existence of a necessary being.

To explain why this is so, I have to develop an idea from the philosophy of mathematics. Until now I've described necessity and contingency as properties of entities. You and I are contingent beings, for example. Now, however, I want to discuss the idea that contingency and necessity are properties of propositions. Propositions, or statements, are true or false. They also have the properties of necessity and contingency.

NECESSARY AND CONTINGENT PROPOSITIONS

Necessity and contingency as properties of propositions are definable in a way that parallels the way this distinction applies to entities like you and me. A proposition is necessarily true if it is true in all possible worlds. A proposition is necessarily false if it is false in all possible worlds. Contingent propositions are true in some possible worlds and false in others. A true statement is contingent if it is true in the actual world, though false in some possible world.

It is easy to cite examples of contingent propositions. Consider the fact that the United States is in North America. This is true, though only contingently so. Consider the following contingent falsehood: The Rocky Mountains are no higher than 6,000 feet above sea level. This is false, but there are possible worlds in which it is true.

MATHEMATICAL TRUTHS

Are there any necessary truths? Many philosophers have held that mathematical truths are necessarily true. Consider the fact that 7 + 3 = 10. It isn't just that this happens to be true in the actual world. There is no possible world in which it is false. The sum of 7 and 3 could not fail to equal 10.

To see why this is at least a plausible view of mathematical propositions, you've got to be clear on what the proposition says. The proposition doesn't say that if you put seven rabbits together with three others, there will be ten bunnies forever after. Rabbits reproduce, but this doesn't contradict any fact of arithmetic. Take another example: If you pour 3 cups of sugar into 7 cups of water, you will not obtain 10 cups of liquid. But, again, this doesn't contradict the fact that 7 + 3 = 10. Arithmetic is unaffected by this result; it just so happens that liquids and solids don't always combine additively.

NAMES DIFFER FROM THE THINGS NAMED

Here is a somewhat more subtle point to notice. The proposition that 7 + 3 = 10 doesn't say anything about the language we use in expressing that fact to ourselves. It is a contingent fact that human beings use the numeral "7" to refer to the number 7. We might have called that number by another name—we might have used the numeral "2" to denote the number 7. Similarly, it is a contingent fact that we use the symbol "=" to represent the relation of equality. So it is a contingent fact that the sentence "7 + 3 = 10" expresses what it does, and it is a contingent fact that that sentence expresses a true proposition rather than a false one. But this doesn't show that 7 + 3 might not have equaled 10. The fact couldn't have been otherwise, even though we might have used our terminology differently.

The distinction I'm drawing here is very obvious in some contexts, though it is easy to lose sight of in others. Few things are more different than a thing and its name. You are a person, but your name isn't. Your name contains a certain number of letters, but you don't. Your name is a linguistic entity, but you aren't.

We mark this difference by using quotation marks. If we want to talk about a thing, we use its name. If we want to talk about the name of a thing, we put the name in quotation marks. So if I want to talk about a big mountain in western North America, I might say that the Rocky Mountains are tall. But if I want to talk about the mountain range's name, I'd say that the phrase "the Rocky Mountains" contains 17 letters.

NUMBERS AREN'T NUMERALS

This is obvious when you think about it. When it comes to mathematics, however, we tend to confuse these ideas. Numerals are names of numbers. But the number and its name are different, just as the phrase "the Rocky Mountains" differs from the mountains themselves. I grant that it is a contingent matter that the sentence

"7 + 3 = 10" expresses a truth. I deny, however, that 7 + 3 might have failed to equal 10. The proposition is necessarily true, even though it is a contingent fact that the sentence expresses the proposition it does. Just as your name is part of a language, though you aren't, so the sentence "7 + 3 = 10" is part of a language, though the proposition it expresses isn't.

SETS

So what does the proposition that 7 + 3 = 10 actually mean? It doesn't say anything about rabbit reproduction or about the way sugar and water combine when poured together. Nor does it say anything about the language we use. You can think of this statement as describing a basic property of mathematical objects called *sets.*

First, I need a bit of terminology. The *union* of two sets is the set that includes all the members of the first and all the members of the second, and no others. A set is any collection of objects. So the set {Groucho Marx, Napoleon, the Eiffel Tower} is a three-member set, and the set {Napoleon, the Rocky Mountains, the French Revolution} is also a three-member set. The items in a set needn't have anything special in common with each other. There are sets of similar things, but there also are sets of dissimilar things. The union of the two sets just mentioned is this set: {Groucho Marx, Napoleon, the Eiffel Tower, the Rocky Mountains, the French Revolution}.

The proposition that 7 + 3 = 10 can be understood to say the following: Consider any two sets where the first contains exactly seven objects and the second contains exactly three. If these sets have no common members, then the union of these two sets will have precisely ten members. Notice that the arithmetic proposition doesn't say that the Rocky Mountains exist, or that there are at least seven material things in the universe. The arithmetic fact doesn't rule out the possibility that the world contains no material objects at all. It simply describes a basic property of the operation of set-theoretic union.

NECESSITY AND CERTAINTY ARE DIFFERENT

I need to add a final clarification of the thesis that mathematical truths are necessary. In saying that it is a necessary truth that 7 + 3 = 10, I'm not saying that I know that the proposition is true with absolute certainty. Nor am I saying that I'll never change my mind on the question of whether 7 + 3 equals 10.

This is a rather subtle point, because we sometimes express our lack of certainty about a proposition by saying that maybe the proposition isn't so. If I think there will be rain tomorrow, but entertain some small doubt that this will be so, I may express this by saying, "Possibly it won't rain tomorrow." But my certainty or uncertainty is a fact about *me.* It is a fact in the *subjective* realm. I want to distinguish this question from the question of whether a given proposition has the property of necessity or contingency. The latter question doesn't have anything special to do with me. Necessity and contingency are *objective.*

Here is a nonmathematical example that may make this clearer. Could there be perpetual motion machines? A machine of this sort wouldn't require any energy input to keep running, but it would provide a constant output of energy. Scientists for many hundreds of years tried to design such a machine. They always failed. Finally, in the nineteenth century, physicists working in the area called thermodynamics proved that perpetual motion machines are impossible.

Consider the proposition that there are or will be perpetual motion machines. This is, we now believe, a falsehood. Is it a necessary falsehood? I would say yes. It isn't just that no one will bother to build one; the point is that it is impossible to build one. This, at least, is what science tells us.

Now I'll ask an entirely separate question: Am I absolutely certain that no such machine will ever exist? I guess I'm not absolutely certain. After all, science has been wrong before, and so maybe it now is wrong when it says that such machines are impossible. Opinion has changed through history on the question of whether such machines could be built. Before the nineteenth century, many serious scientists thought such machines are possible. Later, opinion changed. Maybe it will change again.

So the certainty that a single scientist or a community of scientists may have about the issue may change. But there is something that doesn't change. Either it is possible for such things to exist or it isn't. This should convince you that *certainty* and *necessity* are different. People may change their degree of certainty about a proposition; they may even think at one time that the proposition is true, but later on think that it is false. However, the proposition itself doesn't change from true to false. Nor does a proposition cease to be necessary just because people stop believing it. So, to say that arithmetic truths are necessary isn't to say that people are certain about arithmetic. Nor is it to say that people have never changed their minds about arithmetic propositions. Again, these are questions about our attitudes toward the propositions. But whether a proposition is necessary or not has nothing to do with our attitudes. In this respect, necessity is like truth—both are objective issues, independent of what people happen to believe.

To sum up, I've described a prominent view in the philosophy of mathematics. It holds that arithmetic truths are necessary. I haven't provided that view with a complete defense, but I hope you can see what it asserts. I also hope it is at least somewhat plausible, even if it is not entirely convincing.

What has this material about mathematics to do with the conclusion of Aquinas's third proof? Recall that Aquinas reasons that if there is a necessary being, then God must exist. Part of my point in talking about mathematical necessity here is that I want to describe a necessary being that no one would think is God.

NUMBERS ARE NECESSARY BEINGS

Consider the fact that arithmetic includes various existence statements. Besides asserting that $7 + 3 = 10$, arithmetic also asserts that there is a prime number immediately after 10. If, however, arithmetic truths are necessary, then it is a necessary truth that there exists a prime number immediately after 10. To put it bluntly, the philosophy of mathematics I've described holds that the number 11 is a necessary being. It exists in all possible worlds. This follows from the thesis that arithmetic truths are necessary and from the fact that arithmetic contains existence claims.

The conclusion that there is a God doesn't follow from the assertion that not everything is contingent. A philosopher of mathematics might claim that the number 11 is a necessary being, but this wouldn't entail that God exists. This is my last criticism of Aquinas's third argument.

AQUINAS'S FOURTH ARGUMENT: PROPERTIES THAT COME IN DEGREES

I won't spend much time on Aquinas's fourth proof of the existence of God. It is rooted in an Aristotelian view of causality that seems radically implausible now. Here's the argument:

(1) Objects have properties to greater or lesser extents.

(2) If an object has a property to a lesser extent, then there exists some other object that has that property to the maximum possible degree (call this a maximum exemplar of the property).

(3) So, there is an entity that has all properties to the maximum possible degree.

Hence, God exists.

The argument begins with an observation that is obviously true: Some things are more powerful and some less; some are more intelligent and others less, and so on. The second premise, however, seems entirely implausible. The fact that Charlie Chaplin is less than maximally funny doesn't mean that there must exist a maximally funny comedian. And the fact that we are somewhat intelligent, though not perfectly so, doesn't seem to require that there exists a perfectly intelligent being.

The Aristotelian idea that Aquinas is using here is roughly as follows. Aristotle thought that fire is the maximally hot substance. When other objects are hot to some lesser degree, this is because fire is mixed in them to some extent. The property of heat can occur to a less than maximal degree in human beings, for example, only because there exists this substance, fire, which is hot to the maximum possible degree. Fire is the maximum exemplar of heat, from which lesser degrees of heat derive.

There are other problems with the argument. If each property has a maximum exemplar, it doesn't follow that there is an entity that is a maximum exemplar of all properties. Recall the Birthday Fallacy: The fact that everyone has a birthday doesn't imply that there is a single day on which everyone was born. So even if there is a maximum exemplar of intelligence, a maximum exemplar of power, and a maximum exemplar of moral goodness, it doesn't follow that there is a single entity who is all-knowing, all-powerful, and all-good.

Finally, there is the problem of contradiction. If intelligence has its maximum exemplar, then stupidity would have to have its maximum exemplar as well. By Aquinas's argument, this leads us to say that there is a single being who is both maximally intelligent and maximally stupid. But this is impossible. That concludes what I want to say about Aquinas's fourth way.

CRITICIZING AN ARGUMENT VERSUS SHOWING THE ARGUMENT'S CONCLUSION IS FALSE

None of the four arguments I've discussed here is successful. Does this mean there is no God? It means no such thing. There may be other arguments for the existence of God that are convincing. That the four arguments discussed here don't work doesn't mean that no argument will work. Not one word has been said here that shows that atheism is true. All we have seen is that some arguments for theism are flawed. In the next chapter, I'll consider Aquinas's fifth argument for the existence of God. Maybe it will fare better.

Review Questions

1. What objections are there to the first cause argument?
2. What is the Birthday Fallacy? How does it figure in the discussion of Aquinas's arguments?
3. Explain what it means for an object to be necessary or contingent. What is a "possible world"?
4. How are necessity and eternality related? How does this bear on Aquinas's third argument?
5. What is a reductio argument? Give an example.
6. What is the difference between necessity and certainty? What is meant by saying that necessity is "objective"?
7. What would it mean for something to be a first cause without being God? What would it mean for something to necessarily exist without being God?

Problems for Further Thought

1. I formulated Aquinas's proofs by having him talk about objects that exist in "nature" (in "the natural world"). What does "nature" include? Does it include just the things we can see or hear or touch or taste or smell? After all, the selection from Aquinas in the Readings translates Aquinas as talking about the "sensible" world. How would interpreting his argument in this way affect its plausibility? What other interpretations of "nature" make sense in this context?
2. In discussing Aquinas's third proof, I talked about Charlie the atom as an example of a thing that is both eternal and contingent. Could something exist that is both necessary and noneternal? It would exist at *some* time in each possible world, though it would not exist at *all* times in the actual world. Can you give an example of such a thing?
3. I criticized Aquinas's third argument by discussing numbers, which I claimed exist necessarily. Can the argument be reformulated so that this objection no longer applies?
4. I criticized Aquinas's fourth argument by discussing "maximum stupidity." Can Aquinas reply to this objection by claiming that stupidity is just the absence of intelligence?

CHAPTER 5

The Design Argument

There are three main traditional arguments for the existence of God—the cosmological argument, the design argument, and the ontological argument. Aquinas's first, second, and third ways, surveyed in the previous chapter, are instances of the first. The cosmological argument takes different forms, but all cite general features of the whole universe as evidence that there is a God. The second type of traditional argument—the design argument—is the one we'll consider in the present chapter; the ontological arugment will occupy our attention in Chapter 8.

Aquinas's fifth argument for the existence of God is an instance of what has come to be called the Argument from Design. The design argument has a variety of forms, some of which I'll describe. To start things off, here is a formulation that is close to the one Aquinas uses:

(1) Among objects that act for an end, some have minds whereas others do not.

(2) An object that acts for an end, but does not itself have a mind, must have been designed by a being that has a mind.

(3) Hence, there exists a being with a mind who designed all mindless objects that act for an end.

Hence, God exists.

Note as a preliminary point that the transition from (2) to (3) commits the Birthday Fallacy described in Chapter 4. If each mindless object that acts for an end has a designer, it doesn't follow that there is a *single* designer of all the mindless objects that act for an end.

Readings

SAINT THOMAS AQUINAS

Five Ways to Prove That God Exists

In this selection from ***Summa Theologica*** (Part 1, Question 2, Article 3), Saint Thomas Aquinas provides five proofs of the existence of God. In each, an obvious and uncontroversial observation is the starting premise.

The existence of God can be proved in five ways.

The first and more manifest way is the argument from motion. It is certain, and evident to our senses, that in the world some things are in motion. Now whatever is moved is moved by another, for nothing can be moved except it is in potentiality to that towards which it is moved; whereas a thing moves inasmuch as it is in act. For motion is nothing else than the reduction of something from potentiality to actuality. But nothing can be reduced from potentiality to actuality, except by something in a state of actuality. Thus that which is actually hot, as fire, makes wood, which is potentially hot, to be actually hot, and thereby moves and changes it. Now it is not possible that the same thing should be at once in actuality and potentiality in the same respect, but only in different respects. For what is actually hot cannot simultaneously be potentially hot; but it is simultaneously potentially cold. It is therefore impossible that in the same respect and in the same way a thing should be both mover and moved; *i.e.*, that it should move itself. Therefore, whatever is moved must be moved

Saint Thomas Aquinas, "Five Ways to Prove That God Exists," from *The Basic Writings of St. Thomas Aquinas*, edited by Anton C. Pegis (New York: Doubleday & Co., 1955),

by another. If that by which it is moved be itself moved, then this also must needs be moved by another, and that by another again. But this cannot go on to infinity, because then there would be no first mover, and, consequently, no other mover, seeing that subsequent movers move only inasmuch as they are moved by the first mover; as the staff moves only because it is moved by the hand. Therefore it is necessary to arrive at a first mover, moved by no other; and this everyone understands to be God.

The second way is from the nature of efficient cause. In the world of sensible things we find there is an order of efficient causes. There is no case known (neither is it, indeed, possible) in which a thing is found to be the efficient cause of itself; for so it would be prior to itself, which is impossible. Now in efficient causes it is not possible to go on to infinity, because in all efficient causes following in order, the first is the cause of the intermediate cause, and the intermediate is the cause of the ultimate cause, whether the intermediate cause be several, or one only. Now to take away the cause is to take away the effect. Therefore, if there be no first cause among efficient causes, there will be no ultimate, nor any intermediate, cause. But if in efficient causes it is possible to go on to infinity, there will be no first efficient cause, neither will there be an ultimate effect, nor any intermediate efficient causes; all of which is plainly false. Therefore it is necessary to admit a first efficient cause, to which everyone gives the name of God.

The third way is taken from possibility and necessity, and runs thus. We find in nature things that are possible to be and not to be, since they are found to be generated, and to be corrupted, and consequently, it is possible for them to be and not to be. But it is impossible for these always to exist, for that which can not-be at some time is not. Therefore, if everything can not-be, then at one time there was nothing in existence. Now if this were true, even now there would be nothing in existence, because that which does not exist begins to exist only through something already existing. Therefore, if at one time nothing was in existence, it would have been impossible for anything to have begun to exist; and thus even now nothing would be in existence—which is absurd. Therefore, not all beings are merely possible, but there must exist something the existence of which is necessary. But every necessary thing either has its necessity caused by another, or not. Now it is impossible to go on to infinity in necessary things which have their necessity caused by another, as has been already proved in regard to efficient causes. Therefore we cannot but admit the existence of some being having of itself its own necessity, and not receiving it from another, but rather causing in others their necessity. This all men speak of as God.

The fourth way is taken from the gradation to be found in things. Among beings there are some more and some less good, true, noble, and the like. But *more* and *less* are predicated of different things according as they resemble in their different ways something which is the maximum, as a thing is said to be hotter according as it more nearly resembles that which is hottest; so that there is something which is truest, something best, something noblest, and, consequently, something which is most being, for those things that are greatest in truth are greatest in being, as it is written in *Metaphysics* II [a work of Aristotle]. Now the maximum in any genus is the cause of all in that genus, as fire, which is the maximum of heat, is the cause of all hot things, as is said in the same book. Therefore there must also be something which is to all beings the cause of their being, goodness, and every other perfection; and this we call God.

The fifth way is taken from the governance of the world. We see that things which lack knowledge, such as natural bodies, act for an end, and this is evident from their acting always, or nearly always, in the same way, so as to obtain the best result. Hence it is plain that they achieve their end, not fortuitously, but designedly. Now whatever lacks knowledge cannot move towards an end, unless it be directed by some being endowed with knowledge and intelligence; as the arrow is directed by the archer. Therefore some intelligent being exists by whom all natural things are directed to their end; and this being we call God.

WILLIAM PALEY

The Design Argument

In this selection from his ***Natural Theology*** (1836 edition), William Paley elaborates an argument for the existence of God that traces back at least to the last of Aquinas's five ways. His discussion of how one interprets a watch found on a heath is one of the most famous analogies proposed in the history of philosophy. The term "natural theology" means that the author attempts to establish the existence and nature of God by the same methods used in the natural sciences—observation and reasoning.

In crossing a heath, suppose I pitched my foot against a stone, and were asked how the stone came to be there: I might possibly answer, that for any thing I knew to the contrary, it had laid there for ever: nor would it perhaps be very easy to shew the absurdity of this answer. But suppose I had found a *watch* upon the ground, and it should be inquired how the watch happened to be in that place; I should hardly think of the answer which I had before given, that, for any thing I knew, the watch might have always been there. Yet why should not this answer serve for the watch as well as for the stone? why is it not as admissible in the second case, as in the first? For this reason, and for no other, viz, that, when we come to inspect the watch, we perceive (what we could not discover in the stone) that its several parts are framed and put together for a purpose, e.g. that they are so formed and adjusted as to produce motion, and that motion so regulated as to point out the hour of the day; that, if the different parts had been differently shaped from what they are, of a different size from what they are, or placed after any other manner, or in any other order, than that in which they are placed, either no motion at all would have been carried on in the machine, or none that would have answered the use that is now served by it. To reckon up a few of the plainest of these parts, and of their offices, all tending to one result:—We see a cylindrical box containing a coiled elastic spring, which, by its endeavor to relax itself, turns round the box. We next

William Paley, "The Design Argument," from *Natural Theology*, 1836.

observe a flexible chain (artificially wrought for the sake of flexure), communicating the action of the spring from the box to the fusee. We then find a series of wheels, the teeth of which catch in, and apply to each other, conducting the motion from the fusee to the balance, and from the balance to the pointer; and at the same time by the size and shape of those wheels so regulating that motion, as to terminate in causing an index, by an equable and measured progression, to pass over a given space in a given time. We take notice that the wheels are made of brass in order to keep them from rust; the springs of steel, no other metal being so elastic; that over the face of the watch there is placed a glass, a material employed in no other part of the work, but in the room of which, if there had been any other than a transparent substance, the hour could not be seen without opening the case. This mechanism being observed (it requires indeed an examination of the instrument, and perhaps some previous knowledge of the subject, to perceive and understand it; but being once, as we have said, observed and understood), the inference, we think, is inevitable, that the watch must have had a maker; that there must have existed, at some time, and at some place or other, an artificer or artificers, who formed it for the purpose which we find it actually to answer; who comprehended its construction and designed its use.

I. Nor would it, I apprehend, weaken the conclusion that we had never seen a watch made—that we had never known an artist capable of making one—that we were altogether incapable of executing such a piece of workmanship ourselves, or of understanding in what manner it was performed; all this being no more than what is true of some exquisite remains of ancient art, of some lost arts, and, to the generality of mankind, of the more curious productions of modern manufacture. Does one man in a million know how oval frames are turned? Ignorance of this kind exalts our opinion of the unseen and unknown artist's skill, if he be unseen and unknown, but raises no doubt in our minds of the existence and agency of such an artist, at some former time and in some place or other. Nor can I perceive that it varies at all the inference, whether the question arise concerning a human agent or concerning an agent of a different species, or an agent possessing in some respects a different nature.

II. Neither, secondly, would it invalidate our conclusion, that the watch sometimes went wrong, or that it seldom went exactly right. The purpose of the machinery, the design, and the designer might be evident, in whatever way we accounted for the irregularity of the movement, or whether we could account for it or not. It is not necessary that a machine be perfect, in order to show with what design it was made: still less necessary, where the only question is whether it were made with any design at all.

III. Nor, thirdly, would it bring any uncertainty into the argument, if there were a few parts of the watch concerning which we could not discover or had not yet discovered in what manner they conduced to the general effect; or even some parts, concerning which we could not ascertain whether they conduced to that effect in any manner whatever. For, as to the first branch of the case, if by the loss, or disorder, or decay of the parts in question, the movement of the watch were found in fact to be stopped, or disturbed, or retarded, no doubt would remain in