The Norton Introduction to PHILOSOPHY

Second Edition

GIDEON ROSEN
Princeton University

ALEX BYRNE
Massachusetts Institute of Technology

JOSHUA COHEN
Apple University; University of California, Berkeley

ELIZABETH HARMAN
Princeton University

SEANA SHIFFRIN
University of California, Los Angeles

W. W. NORTON & COMPANY, INC.
New York • London

W. W. Norton & Company has been independent since its founding in 1923, when William Warder Norton and Mary D. Herter Norton first published lectures delivered at the People's Institute, the adult education division of New York City's Cooper Union. The firm soon expanded its program beyond the Institute, publishing books by celebrated academics from America and abroad. By midcentury, the two major pillars of Norton's publishing program—trade books and college texts—were firmly established. In the 1950s, the Norton family transferred control of the company to its employees, and today—with a staff of four hundred and a comparable number of trade, college, and professional titles published each year—W. W. Norton & Company stands as the largest and oldest publishing house owned wholly by its employees.

Printed in the United States of America

Editor: Ken Barton
Assistant Editor: Shannon Jilek
Project Editor: Sujin Hong
Managing Editor, College Digital Media: Kim Yi
Production Manager: Eric Pier-Hocking
Media Editor: Erica Wnek
Media Assistant Editor: Ava Bramson
Marketing Manager, Philosophy: Michael Moss
Design Director: Hope Miller Goodell
Permissions Associate: Elizabeth Trammell
Permissions Manager: Megan Schindel
Composition: S4Carlisle Publishing Services
Manufacturing: LSC Communications—Crawfordsville

ISBN 978-0-393-62442-7

W. W. Norton & Company, Inc., 500 Fifth Avenue, New York, NY 10110
wwnorton.com

W. W. Norton & Company Ltd., 15 Carlisle Street, London W1D 3BS
1 2 3 4 5 6 7 8 9 0

A Brief Guide to Logic and Argumentation

When a philosopher tackles a question, her aim is not just to answer it. Her aim is to provide an argument for her answer and so to present her audience with reasons for believing what she believes. When you read a philosophical text, your main job is to identify and assess the author's arguments. When you write a philosophy paper, your main job is to offer arguments of your own. And because philosophy is an especially reflective discipline—every question *about* philosophy is a philosophical question—philosophers have turned their attention to this phenomenon. What is an argument? What is a *good* argument? How can we tell whether an argument is a good one? The aim of this brief guide is to introduce some of the tools that philosophers have developed for answering these questions. But be warned: some of what follows is controversial, and many of the most important questions in this area remain wide open. It may be unsettling to discover that even at this elementary stage, philosophy raises questions that centuries of reflection have not resolved. But that is the nature of the subject, and you might as well get used to it.

1. WHAT IS AN ARGUMENT?

An **argument** is a sequence of statements. The last claim in the sequence is the *conclusion*. This is the claim that the argument seeks to establish or support. An argument will usually include one or more **premises**: statements that are simply asserted without proof in the context of the present argument but which may be supported by arguments given elsewhere. Consider, for example, the following argument for the existence of God:

ARGUMENT A

(1) The Bible says that God exists.

(2) Whatever the Bible says is true.

(3) Therefore, God exists.

Here the premises are (1) and (2), and statement (3) is the conclusion.

Now, anyone who propounds this argument will probably realize that his premises are controversial, so he may seek to defend them by independent arguments. In defense of (2) he may argue:

ARGUMENT B

(4) The Bible has predicted many historical events that have come to pass.

(5) Therefore, whatever the Bible says is true.

These two arguments may be combined:

ARGUMENT C

(6) The Bible has predicted many historical events that have come to pass.

(7) Therefore, whatever the Bible says is true.

(8) The Bible says that God exists.

(9) Therefore, God exists.

Here the premises are (6) and (8). Statement (7) is now an *intermediate conclusion*, supported by premise (6), and the conclusion of the argument as a whole is (9), which is in turn supported by (7) and (8). It can be useful to make all of this explicit by writing the argument out as follows:

ARGUMENT C, ANNOTATED

(6) The Bible has predicted many historical events that have come to pass. [premise]

(7) Therefore, whatever the Bible says is true. [from (6)]

(8) The Bible says that God exists. [premise]

(9) Therefore, God exists. [from (7), (8)]

All of this is trivial when the arguments are simple and neatly packaged. But when you are reading a philosophical text with an eye toward identifying the author's argument, it is extraordinarily important (and often quite difficult) to distinguish the author's premises—the propositions she takes for granted as a starting point—from her conclusions. Why is this important? If a statement is meant as a conclusion, then it is fair to criticize the author if she has failed to give a reason for accepting it. If, however, a statement is a premise, then this sort of criticism would not be fair. Every argument must start somewhere. So you should not object to an argument simply on the ground that the author has not proved her premises. Of course, you can object in other ways. As we will see, it is perfectly fair to reject an argument when its premises are false, implausible, or defective in some other way. The point is rather simply this: since every argument must have premises, *it is not a flaw in an argument that the author has not argued for her premises.*

> *Rules of thumb:* If a sentence begins with "hence" or "therefore" or "so," that is a clue that it functions as a conclusion. If a sentence begins with "Let us assume that . . ." or "It seems perfectly obvious that . . ." or "Only a fool would deny that . . . ," this is a clue that it functions as a premise.

Exercise: Consider the following passage. What are the premises? What is the main conclusion?

> Everyone knows that people are usually responsible for what they do. But you're only responsible for an action if your choice to perform it was a free choice, and a choice is only free if it was not determined in advance. So we must have free will, and that means that some of our choices are not determined in advance.

2. VALIDITY

An argument is **valid** if and only if it is *absolutely impossible* for its premises to be true and its conclusion false. In our examples, argument A is clearly valid. If the premises are true—if the Bible is infallible, and if the Bible says that God exists—then God must certainly exist. There is no possible situation—no possible world—in which the premises of the argument are true and the conclusion false. Argument B, by contrast, is clearly **invalid**. It is easy to imagine a circumstance in which the Bible makes many correct predictions about historical events while remaining fallible on other matters. When an argument is valid, we say that the premises **entail** or *imply* the conclusion, or, equivalently, that the conclusion *follows from* the premises.

This concept of validity is a technical one, and some of its applications may strike you as odd. Consider:

ARGUMENT D
All philosophers are criminals.
All criminals are short.
Therefore, all philosophers are short.

ARGUMENT E
God exists.
Therefore, God exists.

ARGUMENT F
The moon is green.
The moon is not green.
Therefore, God exists.

It is easy to see that argument D is valid. The premises are *false*, but that is irrelevant. They *could* have been true, and any possible circumstance in which they *are* true is one in which the conclusion is also true. Argument E is also valid. Since the premise and the conclusion are identical, it is clearly impossible for the one to be true and the other false. To see that argument F is valid, note that it is obviously impossible for its premises to be true together—the moon cannot be both green and not green! But this means that it is impossible for the premises to be true and the conclusion false, and that is exactly our definition of validity.

As the examples show, a valid argument can be a *lousy* argument. Still, validity is an important property of arguments. Some disciplines—notably, mathematics—insist on valid arguments at every stage. In these areas, a good argument must be a **proof**, and a proof is a valid argument from premises known to be true. Philosophy, like most disciplines, does not insist on proof. Yet philosophers often aspire to produce valid arguments for their conclusions, and there is a good reason for this. Begin by noting that it is always possible to turn an invalid argument, or an argument whose validity is uncertain, into a valid argument by adding premises. Suppose a philosopher offers the following argument:

> ARGUMENT G
> I can imagine existing without my body. (I can imagine my feet slowly and painlessly disappearing, then my knees, then my legs. . . . As my body disappears, I lose all sensation. As my head disappears, everything goes black and silent because my eyes and ears have disappeared, but still I'm thinking about these strange events, and because I'm thinking, I must exist.)
> Therefore, I am not my body.

It may be hard to say whether this is a valid argument, but we can easily turn it into an argument whose validity is beyond dispute:

> ARGUMENT H
> I can imagine existing without my body.
> If I can imagine *X* existing without *Y*, then *X* is not *Y*.
> Therefore, I am not my body.

A philosopher who offers argument G as a proof that human beings are not identical to their bodies probably has argument H in mind. She is probably tacitly *assuming* the premise that is missing in argument G but that H makes explicit. For philosophical purposes, it is often important to make these tacit assumptions explicit so that we can subject them to the bright light of scrutiny. *When you reconstruct the argument implicit in a philosophical text, you should set yourself the task of producing a valid argument for the author's conclusion from the author's stated premises, supplying any missing premises that might be necessary for this purpose, so long as they are premises that the author might have accepted.* If there are many ways to do this, you will find yourself with several competing interpretations of the argument. If there is only one sensible way of doing this (as with argument G), you will have identified the author's tacit assumptions. This is often a valuable step in your effort to assess the argument.

Exercise: Spot the valid argument(s):

(i) If abortion is permissible, infanticide is permissible.
Infanticide is not permissible.
Therefore, abortion is not permissible.

(ii) It is wrong to experiment on a human subject without consent.
Dr. X experimented on Mr. Z.
Mr. Z consented to this experiment.
Therefore, it was not wrong for Dr. X to experiment on Mr. Z.

(iii) I will not survive my death.
My body will survive my death.
Therefore, I am not my body.

(iv) Geoffrey is a giraffe.
If *X* is a giraffe, then *X*'s parents were giraffes.
Therefore, all of Geoffrey's ancestors were giraffes.

Exercise: The following arguments are not valid as they stand. Supply missing premises to make them valid.

(v) Every event has a cause.
No event causes itself.
Therefore, the universe has no beginning in time.

(vi) It is illegal to keep a tiger as a pet in New York City.
Jones lives in New York City.
Therefore, it would be wrong for Jones to keep a tiger as a pet.

(vii) The sun has risen every day for the past 4 billion years.
Therefore, the sun will rise tomorrow.

Check your understanding. Some statements express *necessary truths*: truths that could not possibly have been false under any circumstances. The truths of pure mathematics are the best examples. There is no possible circumstance in which $2 + 3 \neq 5$, so "$2 + 3 = 5$" is a necessary truth. With this in mind, show that an argument whose conclusion is a necessary truth is automatically a valid argument.

3. SOUNDNESS

A valid philosophical argument is a fine thing. But if the premises are false, it cannot be a good argument. Good arguments, after all, provide us with reasons for accepting their conclusions, and an argument with false premises cannot do that. Recall argument D:

ARGUMENT D

(1) All philosophers are criminals.

(2) All criminals are short.

(3) Therefore, all philosophers are short

The argument is perfectly valid, but it obviously fails to establish its conclusion.

This means that when you evaluate a philosophical argument, it is never enough to show that the author's conclusions follow from her premises. You must also ask whether the premises are true. A valid argument with true premises is called a **sound** argument.

Check your understanding: Use the definitions of soundness and validity to show that if an argument is sound, its conclusion must be true.

4. HOW TO RECONSTRUCT AN ARGUMENT: AN EXAMPLE

One of the most important skills a philosopher can acquire is the ability to extract an explicit argument from a dense block of prose. There is no recipe for doing this: it is an art. Here we work through an example to illustrate one way of proceeding.

Assignment: Identify and assess the argument in the following passage.

> We see that things which lack intelligence, 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 not fortuitously, but designedly, do they achieve their end. Now whatever lacks intelligence cannot move towards an end, unless it be directed by some being endowed with knowledge and intelligence, as the arrow is shot to its mark by the archer. Therefore some intelligent beings exist by whom all natural things are directed to their end. (Thomas Aquinas, *Summa Theologica*, part I, question 2, article 3)

Step 1: Identify the Conclusion

When you seen an argument like this, your first job is to identify the main conclusion. Unsurprisingly, this will usually come at the end, though many writers will tell you at the start what the conclusion of the argument is going to be. (This is very helpful to the reader, and you should always do it in your own writing.) In this case, the main conclusion is helpfully marked by an explicit "therefore."

(Main conclusion) Some intelligent beings exist by whom all natural things are directed to their end.

Step 2: Interpret the Conclusion

Now that you have identified the conclusion, your next job is to understand it. This can be difficult, especially when the text is old and the language unfamiliar. What is it for a being to be *intelligent*? What is a *natural* thing? In this case, the most pressing issue is to understand what it means for a natural thing to be "directed towards an end." As the context makes clear, a natural thing is anything that is not a person or an artifact—an animal or a plant, or perhaps a rock. What is it for such a thing to have an *end*? This is in fact a profound question, but to a first approximation, the end of a thing is its purpose or function. The *end* of the heart is to pump blood, the *end* of a worker bee is to supply food for the queen, and so on. The conclusion of the argument, reformulated in more familiar terms, is therefore this:

(Main conclusion, reformulated)	There is an intelligent being that ensures that natural objects perform their functions.

This illustrates a general point: when you analyze an argument, you are not required to employ the author's original words in every case. It is sometimes useful to supply more familiar words and grammatical constructions, provided they represent a plausible interpretation of the author's meaning. In this case, we have replaced Aquinas's talk of "ends" with talk of "functions."

Step 3: Reconstruct the Argument

Your next job is to reconstruct the argument for the main conclusion. What are the premises from which Aquinas argues? You might think that the first sentence states a premise: "We see that things which lack intelligence . . . act for an end." But as we read on, it becomes clear that this is, in fact, an intermediate conclusion. The first sentence, taken as a whole, is itself an argument.

> Unintelligent things always or nearly always act in the same way, so as to achieve the best result. [premise]
> Therefore, unintelligent things perform a function.

This is an interesting argument, but the connection between the premise and the conclusion is obscure. As it stands, the argument is not clearly valid. But we can render it valid by interpolating an unstated premise:

(1) Unintelligent things always or nearly always act in the same way, so as to achieve the best result.

(2) If a thing always or nearly always acts in a certain way, so as to achieve the best result, then that thing performs a function.

(3) Therefore, unintelligent things perform a function.

This shows the value of making unstated premises fully explicit. The unstated premise (2) contains an important idea. The function of the heart is to pump blood.

How do we know? Because hearts almost always pump blood, and this is a benefit to the organism as a whole. In general, when we see a natural thing acting in a way that provides a benefit, we infer that its function (or one of its functions) is to provide that benefit. The second premise makes this assumption explicit.

When we turn to the next sentence, we have a puzzle. “Hence it is plain that not fortuitously, but designedly, do they achieve their end.” This sentence begins with “hence,” so we naturally assume that it is supposed to be a conclusion supported by what precedes it. If we pursue this interpretation, the argument will look like this:

(1)	Unintelligent things always or nearly always act in the same way, so as to achieve the best result.	[premise]
(2)	If a thing always or nearly always acts in a certain way, so as to achieve the best result, then that thing performs a function.	[premise]
(3)	Therefore, unintelligent things perform a function.	[from (1) and (2)]
	Therefore, unintelligent things perform their functions by design (and not by accident).	[from ?]

The puzzle is that nothing in the argument appears to support this new conclusion. Why shouldn’t natural beings perform their functions by accident rather than by design? Nothing in the text speaks to this question, and so it may be unclear whether Aquinas means this to be a new premise or an intermediate conclusion supported by what comes before.

Again, we can interpolate an unstated premise that will render the argument valid. Aquinas apparently finds it obvious that if a thing has a function, it must have been designed to perform that function. If this is right, then the complete argument up to this point runs as follows:

(1)	Unintelligent things always or nearly always act in the same way, so as to achieve the best result.	[premise]
(2)	If a thing always or nearly always acts in a certain way, so as to achieve the best result, then that thing performs a function.	[premise]
(3)	Therefore, unintelligent things perform a function.	[intermediate conclusion, from (1) and (2)]
(4)	If a thing performs a function, it does so by design.	[implicit premise]
(5)	Therefore, unintelligent things perform their functions by design.	[intermediate conclusion, from (3) and (4)]

The remainder of the argument is now straightforward. The next sentence states another premise.

(6) If an unintelligent thing performs a function by design, then there exists an intelligent being that ensures that it performs this function. [premise]

And from this, Aquinas moves directly to his main conclusion:

(7) Therefore, there exists an intelligent being that ensures that natural objects perform their functions. [conclusion, from (5) and (6)]

What just happened? We took a dense philosophical text and we turned it into an explicit argument. Along the way, we did our best to make the author's unspoken premises explicit and to understand what they might mean. The result is a *reconstruction* of the original argument.

Step 4

We are now in a position to assess the argument as we have reconstructed it. We have two questions to ask: Is it valid, and are the premises true?

Taking the second question first, we twenty-first-century philosophers will have doubts about premise (1)—Do *most* natural things really act so as to achieve the "best result"?—and also about premise (6). The heart of an animal performs a function. Must it have been designed by an intelligent being for that purpose? Certainly not; natural selection can do the job even if no intelligence is involved. So the premises of the argument are certainly open to question.

But even if we waive this objection and suppose that the premises are true, there is a further problem. The conclusion (7) claims there is a *single* intelligent being that ensures that natural things perform their functions. But the premises only require that each natural thing be directed toward its end by some intelligent being or other. To see the difference, note that it is one thing to say that every clock has a designer and another to say that there is a single master-designer who is responsible for every clock. This means that we can accept Aquinas's premises and much of his reasoning without accepting his main conclusion. Even if every natural thing was designed by an intelligent being, it does not follow that a single intelligent being designed them all. Verdict: *Aquinas's argument, as we have reconstructed it, is not valid.*

This brings up a very important point. We have given a reasonably careful reconstruction of Aquinas's argument, but despite our best efforts, the argument as we have reconstructed it is clearly *bad*. Now of course no one is perfect: good philosophers sometimes give bad arguments. But when you have produced a reconstruction of an argument by a good philosopher and the result is an argument that is clearly flawed, that is a sign that you may have misunderstood the original argument. The philosophers represented in this collection are all good philosophers, so you should approach their arguments with this in mind: *Before you dismiss an argument on the basis of your reconstruction of it, you should* *be sure that your reconstruction is the most charitable interpretation you can find.* A charitable reconstruction will present the argument in its best light. It may still involve mistakes, but they will not be gross and obvious mistakes. The most convincing way to object to a philosophical argument is to take the time to identify the best possible version of it, and then to show that *this* version of the argument is still no good.

Exercise: Provide a reconstruction of Aquinas's argument that does not commit the logical error mentioned above in the transition from (6) to (7).

5. FORMAL VALIDITY

Consider:

ARGUMENT I
Every number is an abstract object.
Abstract objects are not located in space.
So numbers are not located in space.

This is a concrete argument with a specific subject matter. It is about numbers, spatial location, and so on. But we can abstract from these specific features of the argument in order to focus on its *form*. One way to do this is to replace all of the subject-specific terms in the argument with *schematic letters*, leaving only the logical skeleton of the argument in place. In the case of argument H, this yields the following *schematic argument*.

Every *F* is a *G*.
*G*s are not *H*.
So *F*s are not *H*.

Once we have identified this schematic argument, it is easy to produce other arguments that exhibit the same form but concern an entirely unrelated subject matter. For example:

ARGUMENT J
Every whale is a mammal.
Mammals do not lay eggs.
So whales do not lay eggs.

In this case, it is clear not just that our original argument is valid but that any argument generated from it in this way must be valid. (The second premise in argument J is false, as every platypus knows. But that does not prevent the argument from being valid. If that puzzles you, review the definition of validity.) When an argument is an instance of a scheme all of whose instances are valid, the argument is said to be **formally valid**.

Note: An argument can be valid without being formally valid. Consider:

ARGUMENT K
Every crayon in the box is scarlet.
So every crayon in the box is red.

The underlying form of this argument is:

Every *F* is *G*.
So every *F* is *H*.

And it is obvious that many arguments of this form will not be valid. (*Exercise:* Give an example.) Of course, we can make argument K formally valid by adding the premise, "If a thing is scarlet, then it is red." As we have emphasized, this is always worth doing when you are analyzing a philosophical argument. And yet, the original argument is valid as it stands, since it is absolutely impossible for the premise to be true and the conclusion false.

Formal logic is the study of formally valid arguments. It aims to catalog the vast array of formally valid arguments and to provide general principles for determining whether any given argument has this feature. Formal logic is an intricate, highly developed subject at the intersection of philosophy and mathematics, and it can be extraordinarily useful for the student of philosophy. Here we list some examples of formally valid arguments along with their traditional names. In what follows, the schematic letters *P*, *Q*, and *R* stand for complete declarative sentences. For your amusement, we also include the standard symbolic representations of these forms of inference. Here "→" means "if . . . then"; "~" means "it is not the case that"; and "∨" means "or."

MODUS PONENS

If *P* then *Q*	$P \to Q$
P	P
Q	Q

MODUS TOLLENS

If *P* then *Q*	$P \to Q$
It is not the case that *Q*	$\sim Q$
It is not the case that *P*	$\sim P$

DISJUNCTIVE SYLLOGISM

Either *P* or *Q*	$P \vee Q$
It is not the case that *P*	$\sim P$
Q	Q

HYPOTHETICAL SYLLOGISM

If *P* then *Q*	$P \to Q$
If *Q* then *R*	$Q \to R$
If *P* then *R*	$P \to R$

CONTRAPOSITION

If *P* then *Q*	$P \to Q$
If it is not the case that *Q*, then it is not the case that *P*	$\sim Q \to \sim P$

All of this may seem obvious, but it can sometimes be quite tricky to determine whether an argument is formally valid. Consider:

A person is responsible for a choice only if it is a free choice.
Every human choice is either caused or uncaused.
If a choice is caused, then it is caused either by prior events or by the agent himself.
If a choice is caused by prior events, then it is not free.
If a choice is uncaused, it is not free.
So a choice is free only if it is caused by the agent himself.
But no choice is caused by the agent himself.
So there is no such thing as a free choice.
So no one is ever responsible for his choices.

Is this a valid argument? You could stare at it for a while, and you might find yourself persuaded one way or the other. Or you could take a logic class and learn enough formal logic to settle the matter conclusively once and for all. One of the great advantages of formal logic is that it permits us to *prove* that an argument of this sort is valid by breaking it down into steps, each of which is indisputably an instance of a valid form.

6. A PUZZLE ABOUT FORMAL LOGIC

Apart from its utility as a tool, formal logic is a source of philosophical perplexity in its own right. Imagine a long row of colored squares on the wall in front of you. The left-most square (square 1) is bright red; the right-most square (square 1000) is bright yellow. The squares in between run from red on the left through orange in the middle to yellow on the right. But there are so many of them that they satisfy the following condition:

(1) Square n and square $n + 1$ are indistinguishable by ordinary means.

If you had a measuring device, you might discover that they differ slightly in color, but you can't tell them apart just by looking, no matter how hard you try. (If you don't think this is possible, get out your paint set and play around. It is easy to produce a sequence of colored patches running from red to yellow that satisfies this condition.)

We now note what appears to be an obvious fact:

(2) If two things are indistinguishable by ordinary means, then if one of them is red, so is the other.

If someone shows you a red rose and tells you, "I've got another rose that's indistinguishable from this one, but it's not red," you would know immediately that he was lying. It's built into our concept of *red* that if two objects look just alike to the naked eye in broad daylight, then either both are red or neither is.

From these two premises, it follows by modus ponens that:

(3) If square n is red, then so is square $n + 1$.

But now we're in trouble. For we can reason as follows:

(4) Square 1 is red.	[premise]
(5) If square 1 is red, then square 2 is red.	[3]
(6) So square 2 is red.	[4, 5, modus ponens]
(7) If square 2 is red, then square 3 is red.	[3]
(8) So square 3 is red.	[6, 7, modus ponens]
. . .	
(1002) So square 999 is red.	[1000, 1001, modus ponens]
(1003) If square 999 is red, then square 1000 is red.	[3]
(1004) So square 1000 is red.	[1002, 1003, modus ponens]

But this is nuts. It was built into our description of the situation that square 1000 is not red; it is bright yellow!

What's gone wrong? If you look closely, you will see that this argument has only three premises. Two of them are stipulated as part of our description of the situation: square 1 is red, and adjacent squares are indistinguishable by ordinary means. The other premise is (2), the claim that there cannot be two indistinguishable things, one of which is red, the other not. The argument uses only one rule of inference: modus ponens. And this leaves us with only two responses to the paradox: either (2) is false and there is a sharp cutoff between red and "not red" somewhere in our series or modus ponens is not a valid rule of inference after all. What is the best response? The problem is called the *sorites paradox* (pronounced saw-*rye*-tees), and it remains unsolved.

7. WHAT MAKES AN ARGUMENT GOOD?

We have seen (see section 2 earlier) that valid arguments can be lousy arguments. The same goes for sound arguments. The question of God's existence is the most important question in the philosophy of religion. But it is easy to produce a sound argument that settles it:

ARGUMENTS L AND M

L: God exists.
  Therefore, God exists.

M: God does not exist.
  Therefore, God does not exist.

These arguments are both formally valid, and one of them has true premises. That means that *one of them is sound*. But neither of these arguments is a contribution to philosophy, and neither could possibly provide a reason for believing its conclusion. Why not?

The obvious answer is that these arguments are defective because they are **circular**—their conclusions are included among their premises—and that is certainly a defect. This might tempt us to say that an argument is good if and only if it is sound and noncircular. But this is not quite right. Consider:

ARGUMENT N

God knows when you will die.
Therefore, God exists.

This argument may be sound, and the premise is clearly *different* from the conclusion, so it is not circular. And yet it is perfectly useless for establishing its conclusion. One way to bring this out is to note that anyone who doubts the conclusion will *automatically* doubt the premise. We cannot imagine a reasonable person *coming to believe* that God exists by first believing that God knows when she will die, and then *inferring* the existence of God. If she believes the premise, she must *already* believe the conclusion.

This shows something important. In a good argument, the premises must be credible *independently* of the conclusion. It must be possible for someone who has not already accepted the conclusion to accept the premise first, and to do so reasonably. This point is sometimes put by saying that a good argument must not **beg the question**. Imagine that you are arguing with someone who doubts your conclusion. Now ask: Could this person reasonably accept my premises if he has not already accepted my conclusion? If not, then the argument is bad in this distinctive way.

It is worth stressing, however, that this idea is not completely clear. Suppose you have read about the platypus, but you are not sure that such things exist. (For all you know, the platypus may be extinct like the dodo or legendary like the hippogriff.) A friend may set you straight as follows:

ARGUMENT O
That thing in the bushes is a platypus.
So platypuses exist.

This is a valid argument, and if it is sound—if your friend really is pointing to a platypus—it might give you an excellent reason for accepting its conclusion. Argument O is thus a good argument: it does not beg the question.

Now suppose that you have been impressed by Descartes's famous suggestion that for all you know, there is no external world at all, and in particular that for all you know, you are a disembodied spirit whose experiences are hallucinations produced in your mind by a malicious demon.[1] At this stage, you are in the market for an argument to show that the material world—the world of rocks and trees and houses—really exists. Trying to be helpful, I hold up a rock and say:

ARGUMENT P
This rock in my hand is a material object.
So material objects exist.

Argument P has exactly the same form as argument O. Both are valid, and both may be sound. And yet it has seemed to many (though not to all) that given the context in which it has been presented, argument P begs the question. If you want to prove the existence of the material world to someone who doubts it, you can't just hold up a rock and say "Voilà!" Your interlocutor, after all, will not believe the rock is real.[2]

1. René Descartes, "Meditation I: What Can Be Called into Doubt," in his *Meditations on First Philosophy,* reprinted in Chapter 6 of this anthology.

2. Samuel Johnson (1709–1784) disagreed. As his biographer reports:

> After we came out of the church, we stood talking for some time together of Bishop Berkeley's ingenious sophistry to prove the nonexistence of matter, and that every thing in the universe is merely ideal. I observed, that though we are satisfied his doctrine is not true, it is impossible to refute it. I never shall forget the alacrity with which Johnson answered, striking his foot with mighty force against a large stone, till he rebounded from it—"I refute it *thus*." James Boswell, *Life of Johnson,* ed. G. B. Hill (Oxford University Press, 1935), Vol. 1, p. 471.

What is the difference between these two "proofs"? This is a difficult question. It is often easy to tell in practice when an argument begs the question—when it *presupposes* what it seeks to prove. But it is quite hard to provide a general rule for determining when an argument begs the question in this sense. This is one point at which our understanding of the contrast between good and bad arguments is incomplete.

8. NON-DEMONSTRATIVE ARGUMENTS

So far we have been discussing valid arguments and asking, in effect: What is the difference between a good valid argument and a bad one? We have seen that a good valid argument must be sound, and that it must not beg the question. And there is no doubt that philosophers have often sought to provide arguments of just this sort. But it would be a grave mistake to suppose that every worthwhile argument must fit this description.

Consider:

ARGUMENTS Q, R, S, AND T

Q: Everyone who has drunk hemlock has died soon afterward.
∴ If I drink this hemlock, I will die.

R: Despite years of looking, no one has ever seen a unicorn.
∴ Unicorns do not exist.

S: The cheese in the cupboard is disappearing.
We hear scratching sounds in the cupboard late at night.
There is a suspicious mouse-sized hole in the back of the cupboard.
∴ A mouse has come to live with us.

T: It's normally wrong to kill a person.
The bartender is a person.
∴ It would be wrong to kill the bartender.

By ordinary standards, these are all excellent arguments. If you are trying to give me reason to believe that unicorns don't exist, or that I will die if I drink the hemlock, or that a mouse has infiltrated the kitchen, or that I shouldn't kill the bartender, these arguments ought to do the trick. But of course *these arguments are not valid*. In each case, it is logically possible for the premises to be true and the conclusion false. Unicorns may be very good at hiding. I may be a biological freak immune to hemlock. The evidence in the kitchen may be a hoax cooked up by my roommates as a joke. The bartender might be a dangerous fiend who will destroy the world unless I shoot him, and so on.

Arguments such as Q, R, S, and T are called **non-demonstrative** arguments. (A *demonstration* is a valid proof, and since these arguments are not valid, they are not demonstrations.) A good non-demonstrative argument must have true premises, and it must not beg the question. But how do we distinguish a good non-demonstrative argument from a bad one? We have a developed theory of validity for demonstrative arguments; namely, formal logic. When it comes to non-demonstrative arguments, however, we have nothing comparable. The problem of formulating a general account of good non-demonstrative reasoning is one of the great open problems in philosophy. We cannot solve it here, but we can introduce some terminology that may be helpful.

Some non-demonstrative arguments exhibit a common form.

Inductive arguments take as premises a series of observations that exhibit a pattern, and then conclude that the pattern holds as a general rule. Argument Q is a very simple inductive argument. Its form appears to be this:

> In the past, events of type A have always been followed by events of type B.
> Therefore, in the future, events of type A will be followed by events of type B.

But it would be a mistake to suppose that every argument of this form is a good one. Consider:

> In the past, every time a presidential election has been held in the United States, the winning candidate has been a man.
> Therefore, in the future, every time a presidential election is held in the United States, the winning candidate will be a man.

As of 2017, the premise of this argument is true; but it would be silly to conclude that there will never be a female president on this basis. Philosophers have long hoped that there might be some sort of formal test for distinguishing the good inductive arguments from the bad ones, but that turns out to be impossible. (See Nelson Goodman's "The New Riddle of Induction" reprinted in Chapter 4 of this anthology.) The theory of statistical inference is an attempt to characterize the good inductive inferences in mathematical terms.

Abductive argument—also called **inference to the best explanation**—begins from some collection of settled facts, and then reasons backwards from these facts to the hypothesis that would best explain them. Arguments R and S are abductive arguments. Their general form is roughly this:

> Certain facts are observed. (The cheese is disappearing, etc.; no one has ever found a unicorn despite years of looking.)
> The best explanation for these facts is *H*. (There is a mouse in the kitchen; there are no unicorns.)
> *H* is a good explanation (and not merely the best of a bad lot).
> Therefore, *H* is (probably) true.

Many of the arguments that one finds in the natural sciences are abductive. Whenever the scientist defends a theory about unobserved objects or events by appeal to evidence, the argument takes roughly this form. (Think about the chemist's case for molecules or Darwin's case for evolution.) A theory of abductive argument will tell us what it is for a hypothesis to constitute the *best* explanation of the data, and it will identify the conditions under which it is reasonable to infer the truth of best explanation. This part of the theory of argumentation is even less well developed than the theory of inductive argument and remains an active area of research.

Argument T is neither inductive nor abductive. Indeed, there is no standard name for arguments of this sort. Their general form is roughly this:

Normally, *P*.
∴ *P*.

We know that cats normally have four legs; so if we are told that Felix is a cat, it is reasonable to infer that Felix has four legs—unless, of course, we have special information about Felix that would suggest that he might be an exception. Arguments of this sort may be especially important in ethics. Some writers hold that the general principles of ethics—unlike the laws of physics and mathematics—are not exceptionless rules but, rather, powerful but imperfect generalizations: rules that hold for the most part, but which tolerate exceptions. If that is so—and this is highly controversial—whenever we apply an ethical principle to a case in order to derive a verdict about how to act, our inference is of this nameless non-demonstrative form.

9. SOME GENERAL REMARKS ON ARGUMENTATION IN PHILOSOPHY

In some areas of inquiry—mathematics is the best example—the only good arguments are valid arguments. Suppose I want to argue for Goldbach's Conjecture: Every even number greater than 2 is the sum of two prime numbers.

If I have a lot of time on my hands, I might begin by checking some examples.

$4 = 2 + 2$ ✓
$6 = 3 + 3$ ✓
$8 = 3 + 5$ ✓
$10 = 3 + 7$ ✓

Impressed by the pattern but getting bored, I might program a computer to check some more examples, and if I do I can easily verify that:

(#) Every even number between 2 and 10 billion is the sum of two prime numbers.

And yet it would be a mistake by mathematical standards to treat this as an argument for Goldbach's Conjecture. It is always *possible* that some even number I have not checked provides a counterexample. The inference from (#) to Goldbach's Conjecture is not valid, and in mathematics the only good arguments are valid arguments.

Philosophy grew up with mathematics, and philosophers have sometimes held themselves to a similar standard, insisting that the only good philosophical arguments are (non–question begging) valid arguments from true premises. (Indeed, they have often insisted on valid arguments from *indisputably* true premises.) This remains the gold standard for argument in philosophy. Interesting arguments of this sort are often possible, and when they are possible, they are desirable. When you reconstruct the arguments of the philosophers for the purposes of evaluating them or when you give arguments of your own, it often makes sense to try for arguments of this sort.

And yet it is a mistake to suppose that philosophical arguments are only good when they are valid. As we have noted, the arguments that serve us well in science and in ordinary life—the arguments that persuade us that atoms and molecules exist or that it would be wrong to kill the bartender—are often non-demonstrative in character. There is no good reason to hold philosophy to a higher standard. But, of course, this leaves us in a difficult position, since as we have stressed, there is no accepted account of when a non-demonstrative argument is a good one.