

Also by Keith Lehrer

*Philosophical Problems and Arguments: An Introduction*  
(coauthored with James W. Cornman and George S. Pappas)

*Knowledge*

*Rational Consensus in Science and Society*  
(coauthored with Carl Wagner)

*Knowledge and Skepticism* (coedited with Marjorie Clay)

*Thomas Reid*

*Metamind*

# Theory of Knowledge

Keith Lehrer

UNIVERSITY OF ARIZONA

Westview Press

BOULDER AND SAN FRANCISCO

from foundationalism and externalism. I thought it important to advocate a theory rather than to pretend to a balanced presentation of views. Total impartiality is unattainable and the attempt at it soporific. Students, quite understandably, like to feel that what they read is a quest for the truth. The current book is, and it reads that way. I am an analytic philosopher who thinks a philosopher, like others, should attempt to define his or her key terms. Definitions chain together as a result, and it requires some intellectual effort to proceed from beginning to end. However, I have explained each definition with an example, which should make it possible for a student who has difficulty with definitions to grasp the main arguments nonetheless.

I recommend the book to those who like argument and definition turned by examples. Those who are seeking effortless mastery of philosophical profundities will not find that here, nor, I think, anywhere else. The students who like to match wits with argumentation and definition should find a feast here. I wish them a hearty meal with good appetite and encourage them to reject what they find unsavory. They should determine the reason for the offensiveness and prepare a dish of their own—one more to their own liking. To understand philosophy, one must do philosophy. One must seek the truth to know it. That is my advice and, as it turns out, my theory of knowledge as well.

In closing I wish to thank Marian David, Scott Sturgeon, Vann McGee, Gary Gleb, and Jonathan Kvanvig for their critical reflections; my research assistant Barbara Hannan and my editor Spencer Carr for reflections and editorial work; and Lois Day for assisting me in preparing the manuscript. I owe special thanks to my research assistant Leopold Stubenberg for proofreading and compiling the index. I should also like to express my indebtedness to the John Simon Guggenheim Foundation and the National Science Foundation for supporting my research on Thomas Reid, which greatly influenced the current work, and the National Endowment for the Humanities for sponsoring the Summer Institute in Theory of Knowledge that I directed with Alvin Goldman. It was this institute that, more than any other single factor, was responsible for my writing this book.

Keith Lehrer

# 1

## The Analysis of Knowledge

ALL AGREE THAT KNOWLEDGE is valuable, but agreement about knowledge tends to end there. Philosophers disagree about what knowledge is, about how you get it, and even about whether there is any to be gotten. The question "What is knowledge?" will be the primary subject of this chapter and of this book. Why approach the theory of knowledge by asking this question? *Epistemology*, the theory of knowledge, and *metaphysics*, the theory of reality, have traditionally competed for the primary role in philosophical inquiry. Sometimes epistemology has won, and sometimes metaphysics, depending on the methodological and substantive presuppositions of the philosopher.

The epistemologist asks what we know, the metaphysician what is real. Some philosophers have begun with an account of the nature of reality and then appended a theory of knowledge to account for how we know that reality. Plato, for example, reached the metaphysical conclusion that abstract entities, or forms, such as triangularity or justice, are real and all else is mere appearance. He also held that the real is knowable, and he inquired into how we might know this reality.<sup>1</sup> Aristotle, on the contrary, held that individual substances, such as individual statues or animals, are real, and inquired as to how we might have knowledge, especially general knowledge, concerning these substances.<sup>2</sup> It is hardly surprising that Plato and Aristotle produced vastly different theories of knowledge when they conceived of the objects of knowledge in such different ways. Their common approach, starting with metaphysics, we might refer to as *metaphysical epistemology*. The problem with this approach is that the metaphysical epistemologist uncritically assumes we know the reality posited and only concerns himself with what such knowledge is like.

Other philosophers, most notably René Descartes, turned tables on the metaphysical approach by insisting that we must first decide what we can *know* about what is real and must remain skeptical about what is real until we have discovered what we can know. We might refer to this as *skeptical epistemology*. However, there is also a problem with this approach. When one once enters the den of skepticism, an exit may be difficult to find. Seeking to discover what he knew by following the method of doubting all that he could, Descartes imagined a powerful demon bent on deceiving us and thus found demonic doubt.<sup>3</sup> It remains controversial whether such doubt admits of relief by reason. It seems natural to begin with skepticism with the hope of discovering what we know and what we do not, but if we first pretend to *total* ignorance, we shall find no way to remove it. Moreover, we shall lack even the meager compensation of knowing that we are ignorant, for that too is knowledge.

Are we then trapped between a method that uncritically assumes our knowledge of reality while assigning priority to metaphysics and one that rejects the assumption that we have knowledge and leads to skepticism? Our approach here will be neither skeptical nor metaphysical. We assign priority to neither metaphysics nor epistemology but attempt to provide a systematic and critical account of prior metaphysical and epistemological assumptions. We refer to this as *critical epistemology*. We begin with commonsense and scientific assumptions about what is real and what is known. These convictions constitute our data, perhaps even conflicting data if commonsense and science conflict. The object of philosophical inquiry, of which critical epistemology is a fundamental component, is to account for the data. The account, though, is critical. Sometimes we explain the data and sometimes we explain the data away. For the most part, it behooves a critical epistemologist to construct a theory of knowledge explaining how we know the things we think we do, but, in a few instances, a theory may explain why we think we know when we do not. In order to explain what we do know or why we do not, however, we do well to first ask what knowledge is. Indeed, we must do so in order to evaluate the claims of either the metaphysical dogmatist or the epistemological skeptic. It is to this inquiry that we now turn.

#### What Is Knowledge?

Some have denied that we know what is true or what is false, and they have remained skeptics. Skepticism will have a hearing, but we shall pursue our study as critical epistemologists: We assume people have knowledge. But what sort of knowledge do they have, and

what is knowledge anyway? There are many sorts of knowledge, but only one, the knowledge that something is true, will be our concern. Consider the following sentences:

- I know the way to Lugano.
- I know the expansion of pi to six decimal places.
- I know how to play the guitar.
- I know the city.
- I know John.
- I know about Alphonso and Elicia.
- I know that the neutrino has a rest mass of 0.
- I know that what you say is true.
- I know the sentence 'Some mushrooms are poisonous' is true.

These are but a few samples of different uses of the word 'know' describing different sorts of knowledge.<sup>4</sup> If we are interested in finding out what people have when they have knowledge, we must first sort out the different senses of the word 'know'. Then we may ask our question again, once it has been disambiguated.

In one sense, 'to know' means to have some special form of competence. Thus, to know the guitar or to know the multiplication tables up to ten is to be competent to play the guitar or to recall the products of any two numbers not exceeding ten. If a person is said to know how to do something, it is this competence sense of 'know' that is usually involved. If I say I know the way to Lugano I mean that I have attained the special kind of competence needed either to get to Lugano or to direct someone there. If I say that I know the expansion of pi expanded to six decimal places, I mean that I have the special competence required to recall or to recite the number pi expanded to six decimal places.<sup>5</sup>

Another sense of 'know' is that in which the word means to be acquainted with something or someone. When I say that I know John, I mean that I am acquainted with John. The sentence 'I know the city' is more difficult to disambiguate. It might mean simply that I am acquainted with the city and hence have the acquaintance sense of 'know', or it might mean that I have the special form of competence needed to find my way around the city, geographically and/or socially. I also might mean that I know it in both the competence and acquaintance sense of 'know'. This example illustrates the important fact that the senses of 'know' that we are distinguishing are not exclusive; thus, the term 'know' may be used in more than one of these senses in a single utterance.<sup>6</sup>

The third sense of 'know' is that in which 'to know' means to recognize something as information. If I know that the neutrino has a rest mass

of 0, then I recognize something as information, namely, that the neutrino has a rest mass of 0. The last three sentences on the list all involve this information sense of the word 'know'. It is often affirmed that to know something in the other senses of 'know' entails knowledge in the information sense of 'know'. I must have some information about Lugano if I know the way to Lugano; about the expansion of pi if I know the expansion of pi to six decimal places; about the city if I know the city; about the guitar if I know how to play the guitar, and so forth. Thus, the information sense of the word 'know' is often implicated in the other senses of the word.

In our study, we shall be concerned with knowledge in the information sense. It is precisely this sense that is fundamental to human cognition and required both for theoretical speculation and practical sagacity. To do science, to engage in experimental inquiry and scientific ratiocination, one must be able to tell whether one has received correct information or not to obtain scientific knowledge of the world. Engaging in law or commerce requires the same sort of knowledge. This sort of knowledge goes beyond the mere possession of information. If you tell me something and I believe you, even though I have no idea whether you are a source of truth and correct information about the subject or a propagator of falsehood and deception, I may, if I am fortunate, acquire information when you happen to be informed and honest. This is not, however, knowledge in the sense that concerns us; it is merely the possession of information. Similarly, if I read some gauge or meter and believe the information I receive, though I have no idea whether the instrument is functioning properly, I may thus acquire information, but this is not knowledge. If you doubt this, consider a clock that is not running because it stopped at noon some months ago. As luck would have it, you happen to look at it just at noon and believe that it is noon as a result. You might, as a result, come to believe it is noon when indeed it is, but that is not knowledge. If the clock is in fact running properly, but, again, you have no idea that this is so, you will have received the information from a reliable source; but your ignorance of the reliability of the source prevents you from recognizing that the information is correct, from knowing that it is correct, even though you may believe it to be so. It is information that we recognize to be genuine that yields the characteristically human sort of knowledge that distinguishes us as adult cognizers from machines, other animals, and even our childhood selves.

Some philosophers, choosing to place emphasis on the similarity between ourselves and these other beings, may insist that they have knowledge when they receive information.<sup>7</sup> This is a verbal dispute in which we shall not engage, for it is profitless to do so. We shall remain

content with the observation that our most cherished scientific achievements, the discovery of the double helix, for example, and our most worthy practical attainments, the development of a system of justice, for example, depend on a more significant kind of knowledge. This kind of knowledge rests on our capacity to distinguish truth from error.

### Analysis

To indicate the information sense of the word 'know' as being the one in question is quite different from analysing the kind of knowledge we have picked out. What is an analysis of knowledge? An analysis is always relative to some objective. It does not make any sense simply to demand the analysis of goodness, knowledge, beauty, or truth, without some indication of what purpose such an analysis is supposed to achieve. To demand the analysis of knowledge without specifying further what you hope to accomplish with it is like demanding blueprints without saying what you hope to build. Before asking for such an analysis, we should explain what goals we hope to achieve with it.

First, let us consider the distinction between analysing the meaning of the term 'know' and analysing the kind of knowledge denoted. Many philosophers have been interested in the task of analysing the meaning of the word 'know'.<sup>8</sup> Indeed, many would argue that there is no need for philosophical analysis once we have a satisfactory analysis of the meaning of the term 'know'. This restrictive conception of philosophical analysis is sustained by a dilemma: either a theory of knowledge is a theory about the meaning of the word 'know' and semantically related epistemic terms, or it is a theory about how people come to know what they do. The latter is not part of philosophy at all, but rather that part of psychology called learning theory. It follows that if a theory of knowledge is part of philosophy, then it is a theory of knowledge about the meaning of the word 'know'. That is the argument, and it is one that would reduce the theory of knowledge to a theory of semantics.

It is not difficult to slip between the horns of the dilemma. A theory of knowledge need not be a theory about the meaning of epistemic words any more than it need be a theory about how people come to know what they do. Instead, it may be one explaining what conditions must be satisfied and how they may be satisfied in order for a person to know something. When we specify those conditions and explain how they are satisfied, then we shall have a theory of knowledge. An analogy should be helpful at this point. Suppose a person says that there are only two kinds of theories about physical mass. Either a theory of matter is a theory about the meaning of 'mass' and semantically related physical terms, or it is a theory about how something comes to have mass. This

dichotomy would be rejected on the grounds that it leaves out the critical question of what mass is, or, to put it another way, it leaves out the question of what condition must be satisfied for something to have a given mass. A theoretician in physics might be concerned with precisely the question of what conditions are necessary and sufficient for an object to have mass, or more precisely, to have a mass of  $n$ . Similarly, a philosopher might be concerned with precisely the question of what conditions are necessary and sufficient for a person to have knowledge, or more precisely, to know that  $p$ .

Some philosophers have questioned whether it is possible to give necessary and sufficient conditions for knowledge, but the finest monuments of scientific achievement mark the refutation of claims of impossibility. Obviously, a necessary and sufficient condition for the application of the expression 'S knows that  $p$ ' is precisely the condition of S knowing that  $p$ . This could be made less trivial with little difficulty. The objection to the idea that a philosopher can discover necessary and sufficient conditions for knowledge may rest on the confused idea that a set of conditions necessary and sufficient for the application of a term constitutes a kind of recipe for applying terms which would enable us to decide quite mechanically whether the term applies in each instance. However, we may, without taking any position on the question of whether such a recipe can be found for applying the term 'know', state flatly that this is not the purpose of our theory of knowledge or the analysis of knowledge incorporated therein. Our interests lie elsewhere.

#### The Form and Objectives of an Analysis of Knowledge

We shall then approach the question "What is knowledge?" with the objectives of formulating necessary and sufficient conditions for a person having knowledge (in the information sense of the term 'know') and of explaining how those conditions may be satisfied. Our project is contiguous with scientific investigations having analogous objectives. Our conception of analysis is indebted to both Carnap and Quine.<sup>9</sup> Carnap proposed that philosophy should aim at *explication*. This is a kind of analysis aiming at the generation of philosophically and scientifically useful concepts. More specifically, explication aims at producing concepts useful for articulating laws and theories. For example, the explication of 'fish' so as to exclude whales from the class of fish generates a scientifically useful concept for the purpose of formulating laws. One such law is that fish are cold-blooded, to which whales would constitute a counterinstance if whales were included in the class of fish. When, however, we take this purpose of explication seriously and adopt the strategy of providing analyses of this sort in philosophy, then, as

Quine argued, there can be no clear boundary between philosophy and science. Our reasoning is that it is surely the purpose of science as well as philosophy to provide concepts to facilitate the formulation of laws and theories.

Thus, we contend that the distinction between philosophy and theoretical science is a bogus distinction, whether viewed historically or systemically.<sup>10</sup> Historically, it is clear that the special sciences break off from philosophy when some theory emerges that deals with a circumscribed subject in a precise and satisfactory manner. Philosophy remains the residual pot of unsolved intellectual problems. To date, theories of knowledge have remained in the pot. We do not claim that the current study or other recent research has brought us to the point where the theory of knowledge should be poured out into a special science, but we hope that we are approaching closer to that goal than some suspect and others fear.

A formulation of an analysis of knowledge may be expressed by an equivalence. Again, the analogy with mass is helpful. An analysis of mass may be given in an equivalence of the following form:

$O$  has a mass of  $n$  if and only if . . .

where the blank to the right of the equivalence is filled with a sentence describing a set of necessary and sufficient conditions. Similarly, an analysis of knowledge may be given in an equivalence of the following form:

$S$  knows that  $p$  if and only if . . .

where the blank to the right of the equivalence is filled with a sentence describing a set of necessary and sufficient conditions.

When considering candidates for such sets of conditions, we shall ask whether there is any counterexample to the proposed analysis. What is a counterexample? First of all, any experiment of fact or thought which would falsify the resulting equivalence is a counterexample. To say that there is no experiment of thought to falsify the equivalence means that we can think of no logical possibility that is consistent with other postulates of the theory under consideration which would yield the result that one side of the equivalence is satisfied and the other is not. We shall begin by considering any logically possible case as a potential counterexample to a theory of knowledge. We may decide eventually, however, that some examples, though logically possible, are so remote in terms of real possibility that they do not constitute realistic objections to an analysis of actual human knowledge.

In addition to being immune from counterexamples, such an equivalence will be a suitable analysis only if it facilitates reaching our epistemic objectives. Thus, though some analyses are definitely mistaken because we can find acceptable counterexamples, there are other equivalences which fail to constitute satisfactory analyses simply because they are unenlightening. To say that a person knows that  $p$  if and only if it is known to the person that  $p$ , though this is immune from counterexamples, would completely fail to explain or inform. The explanatory role of an analysis is of fundamental importance and must be appealed to in support of an analysis.

It is important, therefore, to consider at the outset what sort of enlightenment one is seeking, what one is attempting to explain by means of an analysis. We shall be concerned with an analysis that will be useful for explaining how people know that the input (the reports and representations) they receive from other people and their own senses is correct information rather than error and misinformation. A person may receive a representation that  $p$  as input without knowing that the representation is correct and, therefore, without knowing that  $p$ . Suppose, for example, that some person unknown to me tells me all the perch in the Genesee River will be killed by a pollutant that has raised the temperature of the water two degrees. I might believe what I am told, being gullible, but I do not know whether my informant knows whereof she speaks. Consequently, I do not know the perch will die. My informant may be knowledgeable. I may possess accurate information as a result of believing what I was told, but I do not know that the report is correct. Similarly, if I possess some information in memory but no longer know whether it is correct information, whether it is something I accurately remember or just something I imagine, I am again ignorant of the matter. If, on the other hand, I know that the information I possess is correct, then I have knowledge in the requisite sense.

One test of whether I know that the information I possess is correct is whether I can answer the question of how I know that the information is correct or how I would justify claiming to know. Such questions and the answers provided are the basis for critical discussion and rational confrontation in scientific inquiry and everyday life. The replies to such queries show us whether or not the conditions for knowledge have been satisfied. If a person claims to know something, how well she answers the question "How do you know?" will determine whether we accept her claim. Consequently, our analysis of knowledge should explain how a person knows that her information is correct and how her knowledge claims are justified.

The foregoing remarks indicate why we shall not be concerned with the sort of knowledge attributed to animals, small children, and simple

machines that store information, such as telephones that store telephone numbers. Such animals, children, or machines may possess information and even communicate it to others, but they do not know that the information they possess is correct. They lack any conception of the distinction between veracity and correct information, on the one hand, and deception and misinformation, on the other. Any child, animal, or machine that not only possesses information but knows whether the information is correct is, of course, a candidate for being a knowing subject. In those cases in which such knowledge is lacking, however, we shall assume ignorance in the information sense of knowledge under investigation here.

### The Analysis of Knowledge

With these preliminary remarks to guide us, we shall now offer an analysis of knowledge. Each condition proposed will be the subject of subsequent chapters. Moreover, in the case of some controversial conditions, we shall not undertake a detailed defense in the present chapter. Our intention here is only to provide the analysis with some intuitive justification which will subsequently be developed and defended.

### A Truth Condition

The first condition of knowledge is that of *truth*. If I know that the next person to be elected President of the United States will have assets of at least one million dollars, then it must be true that the next President will have assets of at least one million dollars. Moreover, if the next person to be elected President will, in fact, not have assets of at least a million dollars, then I do not know the next President will have assets of at least a million dollars. If I claim to know, my knowledge claim is incorrect. I did not know what I said I did. Thus, we shall accept the following conditionals:

(iT) if  $S$  knows that  $p$ , then it is true that  $p$

and

(iT') If  $S$  knows that  $p$ , then  $p$ .

The two conditionals are equivalent for all those cases in which instances of the following principle, which articulates the absolute theory of truth to be discussed in the next chapter, are necessarily true:

(AT) It is true that  $p$  if and only if  $p$ .

It is true that the U.S. has a president if and only if the U.S. has a president, and this is necessarily true. The equivalence of the conditionals

If Lehrer knows that the U.S. has a president, then the U.S. has a president

and

If Lehrer knows that the U.S. has a president, then it is true that the U.S. has a president

is a result of the necessary truth of

It is true that the U.S. has a president if and only if the U.S. has a president.

We shall find in the next chapter, however, that in spite of the innocent and even trivial appearance of (AT), the absolute theory of truth, it leads to paradox in some instances.

#### *An Acceptance Condition*

The second condition of knowledge is *acceptance*. If I deceitfully claim to know that Jan and Jay married on 31 December 1969, when I do not accept it, then I do not know Jan and Jay married on that date even if they were married then. If I do not accept that  $p$ , then I do not know that  $p$ . Thus, the following conditional expresses a condition of knowledge:

(iA) If  $S$  knows that  $p$ , then  $S$  accepts that  $p$ .

A more familiar and quite similar condition would require belief as a condition of knowledge as follows:

(iB) If  $S$  knows that  $p$ , then  $S$  believes that  $p$ .

These two conditions would be equivalent if the following equivalence were necessarily true:

(AB)  $S$  accepts that  $p$  if and only if  $S$  believes that  $p$ .

Principle (AB) is not true, however. There is a special kind of acceptance requisite to knowledge. It is accepting something for the purpose of attaining truth and avoiding error with respect to the very thing one accepts. More precisely, the purpose is to accept that  $p$  if and only if  $p$ . Sometimes we believe things that we do not accept for this epistemic purpose. We may believe something for the sake of felicity rather than from a regard for truth. We may believe that a loved one is safe because of the pleasure of so believing, though there is no evidence to justify accepting this out of regard for truth, indeed, even when there is evidence against it. So, there are cases in which we do not accept in the appropriate way what we believe. It is the acceptance of something in the quest for truth that is the required condition of knowledge.

Some philosophers have insisted that a person may know something is true even though she lacks conviction of its truth. Others, in diametric opposition, have contended that a person only knows that something is true when she is sure, or certain, of the truth of what she believes. Thus, some philosophers have denied condition (iB) on the grounds that a person may know something to be true that she does not believe at all,<sup>11</sup> and others have maintained that for a person to know something to be true she must believe it to be true with considerable certainty.<sup>12</sup> Our proposal is that acceptance rather than belief, condition (iA) rather than (iB), is what is needed. A person need not have a strong feeling that something is true in order to know that it is. What is required is acceptance of the appropriate kind, acceptance in the interest of obtaining a truth and avoiding an error in what one accepts.

We may, however, consider the appropriate kind of acceptance to be a kind of belief, provided we do not assume that all kinds of belief are the requisite sort of acceptance. Hence, we might adopt

(A) If  $S$  accepts that  $p$ , then  $S$  believes that  $p$ ,

provided that we reject

(B) If  $S$  believes that  $p$ , then  $S$  accepts that  $p$ .

We gain some continuity with tradition as well as some expository simplification by considering acceptance to be a special kind of belief. We may, consequently, speak of belief as a condition of knowledge for the sake of tradition, but we shall recall that it is a special kind of belief—acceptance aimed at truth—that is required and introduce the terms “accept” and “acceptance” when precision is needed.

### A Justification Condition

Accepting something that is true does not suffice for knowledge. If I accept something without evidence or justification, that my wife has exactly fourteen dollars in her purse, for example, and, as luck would have it, this turns out to be right, I fall short of *knowing* that what I have accepted is true. Thus, we require a third condition affirming the need for *justification*. While we allowed that a person need not be completely certain of *p* in order to know that *p*, we shall insist that he be justified, indeed, completely justified in his acceptance of *p* in order to be said to know that *p*.

The reason for requiring that a person be *completely* justified rather than simply justified is to indicate that slight justification is not enough. I may be justified in accepting that my secretary is in her office now because she is ordinarily there at this time. Not being there myself, however, I do not know that she is there, for, though justified, I am not completely justified in accepting that she is there. I am unable to exclude the possibility that she is out of the office on an errand, for example, and, in that way, my justification is incomplete. Our condition may be formulated as follows:

(ij) If *S* knows that *p*, then *S* is completely justified in accepting that *p*.<sup>13</sup>

The locution '*S* is completely justified in accepting that *p*' will be used in a somewhat technical way. We offer some clarification of what is meant here, but the analysis of this notion must be left to later chapters. In colloquial usage, a speaker may say that another is completely justified in accepting that *p* because the *speaker* rather than the other person has strong evidence that *p*. There may be no implication that the other has such evidence. For example, if someone says, 'Alexander believes his wife is unfaithful' and I reply, 'He is completely justified', I may be implying only that I have adequate evidence of her infidelity, never mind how I acquired it, without any implication that Alexander has such evidence. Thus, I could expand the previous utterance and say instead, 'He is completely justified as it happens, but he really has no evidence of her infidelity—she is too clever'.

This use of the expression 'completely justified' is not the one intended in (ij). When we say that *S* is completely justified, we shall mean that if his acceptance is based on adequate evidence, then he is completely justified by the evidence he has in accepting that *p*. Thus, that I am completely justified in accepting that *p* by the evidence I have does not by itself warrant my saying that another is completely justified in her

acceptance of *p*. She too must have evidence which completely justifies her acceptance before she is, in the required sense, completely justified in accepting that *p*. The moral of the preceding remarks is that we shall not be enslaved to ordinary thought and speech when we speak of "complete justification" but, for the sake of theoretical advantage, we shall delete unwanted implications and allow expedient expansion within the theory of justification articulated below.

### Theories of Justification

There are three kinds of theories of justification that we shall discuss in detail in subsequent chapters. These theories constitute the heart of a theory of knowledge. The first kind of theory is a *foundation* theory of justification. According to foundationalists, knowledge and justification are based on some sort of foundation, the first premises of justification. These premises provide us with basic beliefs that are justified in themselves, or self-justified beliefs, upon which the justification for all other beliefs rests.<sup>14</sup>

The motives for such a theory are easy to appreciate. If one thinks of justification in terms of an argument for a conclusion, it appears that justification must either continue infinitely from premise to premise, which would be an infinite regress, or argumentation must cycle with some premise being used to justify itself. This would be a circular argument, or some premises must be first premises, for example, basic beliefs justified without appeal to other premises. The latter alternative is the one chosen by the foundation theory.

Basic beliefs constitute the evidence in terms of which all other beliefs are justified according to the foundation theorist. Some empiricist philosophers affirm the existence of basic beliefs concerning perception (*I see something red*, for example), or more cautious beliefs about mere appearance (*I am appeared to in a reddish way*, for example) and maintain that all justification would be impossible without them. They aver that unless there are some basic beliefs to which we may appeal in justification, we shall lack a necessary starting point and fall victim to skepticism. In the absence of basic beliefs the whole edifice of justification would collapse for want of a foundation.

Not all epistemologists agree with this contention. A second kind of theory of justification, a *coherence* theory, denies the need for basic beliefs. Coherentists argue that justification must be distinguished from argumentation and reasoning. For them, there need not be any basic beliefs because all beliefs may be justified by their relation to others by mutual support.<sup>15</sup> The edifice of justification stands because of the



way in which the parts fit together and delicately support each other rather than because they rest on a concrete foundation of basic beliefs.

How can a theory of justification avoid an unceasing regress proceeding from premise to premise without appeal to basic beliefs? First, justification need not proceed until all claims to knowledge employed in the justification are themselves justified. If we consider justification in a social context, the justification of knowledge claims need proceed only as long as some claim to knowledge is disputed. Thus, if we suppose that justification is a response to a query or demand, then there is no reason to suppose that the argument need proceed beyond the point at which agreement is reached. Hence, even if all completely justified beliefs are justified by evidence, not all claims to knowledge employed to defend other such claims need themselves be justified. They need to be justified only when they engender disputation. Just as we avoid endless disputation by finding premises on which we agree, we may avoid a regress of justification without appeal to basic beliefs, says the coherence theorist, because beliefs are completely justified by the way they agree or cohere with a system of beliefs. My perceptual belief that I see something red, for example, is justified because of the way it coheres with a system of beliefs that tells me under what conditions I can tell something red when I see it. It is coherence rather than reasoning or argumentation which yields justification.

This dispute between the foundation theorist and the coherence theorist is joined by a third party, the defender of an *externalist* theory, who disagrees with both parties to the dispute. We need neither basic beliefs nor coherence to obtain knowledge, the externalist contends, but rather the right sort of external connection between belief and reality to obtain knowledge.<sup>16</sup> Causality is one contender for the role of the needed external connection. What makes my belief that I see something red a case of knowledge on such an account is that my belief is *caused* by my seeing some external red object. Such philosophers may even go so far as to deny that justification is necessary for knowledge, contending that only the desired external connection is necessary. We may, however, do the externalist no injury by looking upon the external connection as providing us with a kind of external justification.

The foundation theorist and the coherence theorist may together protest, however, that a person totally ignorant of the external relationship of her belief, the causal history of her belief, for example, will not know that her belief is true unless it is justified by basic beliefs or coheres with a system of beliefs. The externalist will reply that the appropriate external connection requires neither basic beliefs nor coherence to yield knowledge. We leave the dispute unresolved here to become the centerpiece of our inquiry.

We shall eventually argue, however, that complete justification is a matter of coherence within a system of things a person accepts, which is a *subjective* fact about the knower but with some features adapted from the foundation theorist and the externalist. From the former, we shall take the insight that some beliefs are justified without being conclusions of argumentation and, from the latter, we shall incorporate the idea that a system yielding coherence may contain correct representations of how our beliefs are connected to reality. Nevertheless, we shall find that the engine of justification, which pulls the epistemic lever, is something subjective, something a person accepts in the quest for truth.

Most philosophers have thought that knowledge must be based on some objective method for assessing claims of truth or falsity. Some thought the test was that of experience, others of reason, and there have been mixed methodologies as well. All have assumed that acceptance must be checked in some objective manner. They have repudiated with epistemic horror the idea that acceptance of any sort could by itself produce any sort of justification. That a person accepts something for whatever purpose is far too subjective a datum to serve as a solid basis for justification. Even those philosophers who argue that some beliefs are self-justified have sought some principle by means of which we can determine which beliefs are self-justified and which not. They have held, too, that we must somehow transcend the subjectivity of acceptance in order to demarcate the area of justification. This conception has become so ingrained philosophically as to impose itself on commonsense. However, the assumption that there is some objective method for distinguishing the honest coin of justified acceptance from the counterfeit of the unwarranted shall not go unexamined. We shall study in some detail those theories that rest on this assumption, but, to warn the reader fairly in advance, no such theory shall prevail once we have exhibited our mint for epistemic approval.

The theory of justification we shall ultimately defend may strike some as closely aligned with skepticism. We shall examine this charge, but even here it should be noted that our sympathies with the writings of the philosophical skeptics of the past are strong. Too often contemporary writers seek the most effective method for liquidating the skeptic without asking whether his teaching may not be of more importance than his mode of burial. Since the most brilliant philosophers of past and present have been skeptics of one form or another, it would behoove those who study skepticism to consider whether these skeptics have some truth in their grasp. We claim they do. At the heart of the skeptic's position is the insight that there is no exit from the circle of what one accepts from which one can sally forth to find some exquisite tool to measure the

merits of what lies within the circle of subjectivity. Nor is there such a tool, as we shall show, but subjectivity when directed toward truth and away from error can provide the basis of complete justification.

#### A Counterexample

Some philosophers have suggested that the conditions which we have considered necessary for knowledge are jointly sufficient for knowledge as well.<sup>17</sup> This would amount to affirming the following equivalence as an analysis of knowledge:

*S* knows that *p* if and only if it is true that *p*, *S* accepts that *p*, and *S* is completely justified in accepting that *p*.

In short, knowledge is completely justified true acceptance. Nevertheless, this analysis has been forcefully disputed and requires amendment.<sup>18</sup>

Edmund Gettier has presented us with a counterexample to the claim that knowledge is completely justified true acceptance which runs as follows. Suppose a teacher wonders whether any member of her class owns a Ferrari and, moreover, suppose that she has very strong evidence that one student, a Mr. Nogot, owns a Ferrari. Mr. Nogot says he does, drives one, has papers stating he does, and so forth. The teacher has no other evidence that anyone else in her class owns a Ferrari. From the premise that Mr. Nogot owns a Ferrari, she draws the conclusion that at least one person in her class owns a Ferrari. The woman might thus be completely justified in accepting that Mr. Nogot owns a Ferrari.

Now imagine that, in fact, Mr. Nogot, evidence to the contrary notwithstanding, simply does not own the Ferrari. He was out to deceive his teacher and friends to improve his social status. However, another student in the class, a Mr. Havit, does own a Ferrari, though the teacher has no evidence or inkling of this. In that case, the teacher would be quite correct in her belief that at least one person in her class owns a Ferrari, only it would not be Mr. Nogot who she thinks owns one, but Mr. Havit instead. In this case, the teacher would have a completely justified true belief when she accepts that at least one person in her class owns a Ferrari, but she could not be said to know that this is true because it is more due to good fortune than good justification that she is correct.<sup>19</sup>

To put the argument schematically, Gettier argues that a person might be completely justified in accepting that *F* by her evidence, where *F* is some false statement, and deduce *T* from *F*, where *T* is some true statement. Having deduced *T* from *F*, which she was completely justified in accepting, the person would then be completely justified in accepting

that *T*. Assuming that she accepts that *T*, it would follow from the analysis that she knows that *T*. In such a case, the belief that *T* will be true, but the only reason the person has for accepting *T* to be true is the inference of *T* from *F*. Since *F* is false, it is a matter of luck that she is correct in her belief that *T*.<sup>20</sup>

One might be inclined to reply that inference from a false statement can never yield complete justification, but similar examples may be found that do not seem to involve any inference. An example taken from R. M. Chisholm illustrates this. Suppose a man looks into a field and spots what he takes to be a sheep.<sup>21</sup> The object is not too distant and the man knows a sheep when he sees one. In such a case, it would be natural to regard the man as being completely justified in accepting that he sees a sheep in the field without any reasoning at all. Now imagine that the object he takes to be a sheep is not a sheep but a dog. Thus, he does not know that he sees a sheep. Imagine, further, that an object in the deeper distance which he also sees but does not think is a sheep, happens in fact to be a sheep. So it is true that the man sees a sheep and, moreover, accepts and is completely justified in accepting that he sees a sheep. Of course, he still does not know that he sees a sheep because what he takes to be a sheep is not, and the sheep that he sees he does not take to be a sheep.

#### Justification Without Falsity: A Fourth Condition

In the two cases we have described, a person has justified true acceptance but lacks knowledge and in one case does not infer what he thus accepts from any false statement. There is some merit, however, in the idea that the falsity of some statement accounts for the lack of knowledge. Somehow, it is the falsity of the two statements (that Mr. Nogot owns a Ferrari and that what the man takes to be a sheep really is one) which accounts for the problem. It is false that Mr. Nogot owns a Ferrari, and it is also false that what the man takes to be a sheep is really a sheep (because it is a dog). We may say that in the first case the teacher's justification for her belief that at least one person in her class owns a Ferrari depends on the false statement that Mr. Nogot owns a Ferrari, and in the second case that the man's justification for his belief that there is a sheep in the field depends on the false statement that what he takes to be a sheep is really a sheep.

We shall explore the kind of dependence involved subsequently, but here we may notice that the teacher would be unable to justify completely her acceptance that there is a Ferrari owner among her students were she to concede the falsity of the statement that Mr. Nogot owns a Ferrari. Similarly, the man would be unable to justify completely his acceptance

that there is a sheep in the field were he to concede the falsity of the statement that what he takes to be a sheep really is a sheep.

To render our analysis impervious to such counterexamples, we must add the condition that the complete justification that a person has for what she accepts must not depend on any false statement—whether or not it is a premise in inference. We may thus add the following condition to our analysis:

(iD) If  $S$  knows that  $p$ , then  $S$  is completely justified in accepting that  $p$  in some way that does not depend on any false statement.<sup>22</sup>

#### A Final Analysis of Knowledge

The preceding condition enables us to complete our preliminary analysis of knowledge as follows:

(AK)  $S$  knows that  $p$  if and only if (i) it is true that  $p$ , (ii)  $S$  accepts that  $p$ , (iii)  $S$  is completely justified in accepting that  $p$ , and (iv)  $S$  is completely justified in accepting  $p$  in some way that does not depend on any false statement.

Our next task is to examine each of these conditions of knowledge in order to formulate a theory of knowledge explaining how and why claims to knowledge are justified. We begin in the next chapter with an account of truth and acceptance and then proceed to consider theories of justification. The discussion of such theories will lead us to an account that brings central features of the various theories under the umbrella of a coherence theory. The correct theory of knowledge must provide the correct blend of subjective acceptance and truth in what is accepted, the right match between mind and reality. A match between mind and world sufficient to yield knowledge rests on coherence with a system of things we accept, our acceptance system, which must include an account, undefeated by error, about how we may succeed in our quest for truth. When we have such a theory before us, we shall return, at the end, to the speculations of skeptical and metaphysical epistemologists supplied with the scale of knowledge to weigh their claims.

#### Introduction to the Literature

There are a number of good introductions to the theory of knowledge. Perhaps the best general collection of essays pertaining to both the classical and contemporary literature is *Human Knowledge*, edited by Paul K. Moser and Arnold Vander Nat. The best collection of

contemporary articles is *Essays on Knowledge and Justification*, edited by George S. Pappas and Marshall Swain. Two splendid and readable traditional introductions are *The Problems of Philosophy* by Bertrand Russell and *The Problem of Knowledge* by Alfred J. Ayer. There are some excellent recent textbooks written by single authors. The best are *Contemporary Theories of Knowledge* by John L. Pollock, *Belief, Justification, and Knowledge* by Robert Audi, and *Theory of Knowledge*, 3rd ed., by Roderick Chisholm.