

# PART III

# THEORY OF KNOWLEDGE

---

## LECTURE 12

---

### *What Is Knowledge?*

#### EPISTEMOLOGY

In everyday life, in science, and in philosophy as well, we talk of “knowing” things. We also say that some beliefs are “strongly supported by evidence”; we say that they are “justified” or “well-confirmed.” We not only describe *ourselves* with such terms; we also apply them to *others*.

*Epistemology* is the part of philosophy that tries to understand such concepts. Epistemologists try to evaluate the commonsense idea that we (often, if not always) have knowledge and that we are (often, if not always) rationally justified in the beliefs we have. Some philosophers have tried to defend these commonsense ideas with philosophical argumentation. Others have developed a philosophical position that involves denying these commonsense ideas. A philosopher who claims that we don’t have knowledge, or that our beliefs aren’t rationally justified, is defending some version of *philosophical skepticism*.

In this lecture, I’ll begin with some remarks about the problem of knowledge. In the next lecture, I’ll examine the views of the seventeenth-century French philosopher René Descartes. Descartes tried to show that we really do possess knowledge of the world; Descartes tried to refute the skeptic. After evaluating Descartes’s views about knowledge, I’ll turn to an alternative way of thinking about knowledge, *the Reliability Theory of Knowledge* (Lecture 14).

After this discussion of the problem of knowledge, I’ll examine the idea of rational justification. The eighteenth-century Scottish philosopher David Hume, whose views on the Argument from Design we considered in Lecture 5, argued that the beliefs we have that are based on induction aren’t rationally justified. Hume was a skeptic about induction. We’ll consider his argument for this philosophical position and also the attempts some philosophers have made to show that Hume’s startling thesis is mistaken.

The last lecture in this part of the book introduces ideas about probability that bear on Hume's argument and also have implications for practical problems that arise in everyday life (for example, medical diagnosis).

So let's get started with the problem of knowledge. Before we can ask whether we know anything, we have to get clear on what knowledge is. To focus ideas, I want to distinguish three different ways we talk about knowledge. Only one of these will be our concern in what follows.

### THREE KINDS OF KNOWLEDGE

Consider the differences that separate following three statements, each concerning an individual, whom I'll call *S* (the subject):

- (1) *S* knows how to ride a bicycle.
- (2) *S* knows the President of the United States.
- (3) *S* knows that the Rockies are in North America.

Right now I'm not interested in saying which of these statements is true. The point is that they involve different kinds of knowledge.

The kind of knowledge described in (3) I'll call *propositional knowledge*. Notice that the object of the verb in (3) is a proposition—something that is either true or false. There is a proposition—that the Rockies are in North America—and (3) asserts that *S* knows that that proposition is true.

Statements (1) and (2) don't have this structure. The object of the verb in (2) isn't a proposition, but a person. A similar kind of knowledge would be involved if I said that *S* knows Chicago. Statement (2) says that *S* is related to an object—a person, place, or thing—so I'll say that (2) describes an instance of *object knowledge*.

Is there a connection between object knowledge and propositional knowledge? Maybe to know the President of the United States, you must know some propositions that are about him. But which propositions? To know the President, do you have to know what state he comes from? This doesn't seem essential. And the same holds true for each other fact about the man: There doesn't seem to be any particular proposition you've got to know for you to know him.

There is another aspect of the idea of object knowledge, one that is rather curious. Suppose I've read lots of books about the President. I know as many propositions about him as you might wish. Still, it won't be true that I know him, because I've never met him. Knowing people seems to require some sort of direct acquaintance. But it is hard to say exactly what is needed here. If I once was introduced to the President at some large party, that wouldn't be enough for me to say that I "know" him. It isn't just direct acquaintance, but something more. I won't try to describe this further. I'll merely conclude that propositional knowledge, no matter how voluminous, isn't sufficient for object knowledge.

I turn next to the kind of knowledge described in statement (1). I'll call this *know-how* knowledge. What does it mean to know how to do something? I think this idea

I'm familiar with the  
 contact not lang

has little connection with propositional knowledge. My son Aaron knew how to ride a bike when he was five years old. This means he had certain abilities—he knew how to keep his balance, how to pedal, and so on. If you asked a physicist to describe what Aaron was doing that allowed him to ride the bike, the physicist could write out a set of propositions. There would be facts about gravity, forward momentum, and balancing of forces. But Aaron wasn't a physicist at age five. He didn't know the propositions that the physicist specifies. Aaron *obeyed* the physical principles that the physicist describes—his behavior conformed to what they say. But he didn't do this by learning the propositions in question. Aaron had know-how knowledge, but little in the way of propositional knowledge.

### Necessary and Sufficient Conditions ✓

Consider the following suggested definition of what a bachelor is:

For any *S*, *S* is a bachelor if and only if

- (1) *S* is an adult.
- (2) *S* is male.
- (3) *S* is unmarried.

Suff.    iff

I'm not claiming that this definition precisely captures what "bachelor" means in ordinary English. Rather, I want to use it as an example of a proposed definition.

The definition is a generalization. It concerns any individual you care to consider. The definition makes two claims: The first is that *IF* the individual has characteristics (1), (2), and (3), then the individual is a bachelor. In other words, (1), (2), and (3) are together *sufficient* for being a bachelor. The second claim is that *IF* the individual is a bachelor, then the individual has all three characteristics. In other words, (1), (2), and (3) are each *necessary* for being a bachelor.

We can define what a necessary condition is and what a sufficient condition is as follows:

"*X* is a necessary condition for *Y*" means that if *Y* is true, then *X* is true.

"*X* is a sufficient condition for *Y*" means that if *X* is true, then *Y* is true.

What does the expression "if and only if" mean in the above proposed definition of bachelorhood? It means that the conditions listed are both necessary and sufficient. A good definition will specify necessary and sufficient conditions for the concept one wishes to define.

This means that there can be two sorts of defect in a proposed definition. A definition can fail to provide conditions that are sufficient. It also can provide conditions that aren't necessary. And, of course, it can fall down on both counts.

Which sorts of defects are present in the following suggested definitions?

*S* is a bachelor if and only if *S* is male.

*S* is a bachelor if and only if *S* is an unmarried human adult male who is tall and lives in Ohio.

The first definition is said to be "too broad"; it admits too much. The second is "too narrow"; it admits too little.

Conversely, it is possible for a physicist to have detailed knowledge of the physical principles that describe successful bike riding and yet not know how to ride a bike. The physicist may lack the ability to perform the behaviors, but not because there is some proposition he or she fails to grasp.

I conclude that propositional knowledge is neither necessary nor sufficient for knowing how to perform some task. The two concepts of knowledge are quite separate.

The subject of this and the following two lectures will be *propositional knowledge*, not *object knowledge* or *know-how knowledge*. The goal is to understand what propositional knowledge is. That is, we want to answer the following question: What are the necessary and sufficient conditions for it to be true that *S* knows that *p*, where *p* is some proposition—for example, the proposition that the Rockies are in North America?

## TWO REQUIREMENTS FOR KNOWLEDGE: BELIEF AND TRUTH

Two ideas that form part of the concept of knowledge should be noted at the outset. First, if *S* knows that *p*, then *S* must believe that *p*. Second, if *S* knows that *p*, then *p* must be true. Knowledge requires both *belief* and *truth*.

I won't try to argue for the first of these requirements, but will just assume it is correct. The second does require some explanation, however. People sometimes say they know things that, in fact, turn out to be false. But this isn't a case of knowing things that are untrue, but of people *thinking* they know things that happen to be untrue.

Knowledge has an objective and a subjective side. You should remember this pair of concepts from Lecture 1. A fact is objective if its truth doesn't depend on the way anyone's mind is. It is an objective fact that the Rocky Mountains are more than 10,000 feet above sea level. A fact is subjective, on the other hand, if it isn't objective. The most obvious example of a subjective fact is a description of what's going on in someone's mind.

Whether someone believes the Rockies are more than 10,000 feet above sea level is a subjective matter; but whether the mountains really are that high is an objective matter. Knowledge requires both an objective and a subjective element. For *S* to know that *p*, *p* must be true and the subject, *S*, must believe that *p* is true.

I've just cited two necessary conditions for knowledge: Knowledge requires belief, and knowledge requires truth. Is that it? That is, are these two conditions not just separately necessary but also jointly sufficient? Is true belief enough for knowledge?

## PLATO: TRUE BELIEF ISN'T SUFFICIENT FOR KNOWLEDGE

In the dialogue called *The Theaetetus*, the Greek philosopher Plato (c. 430–345 B.C.), who was Aristotle's teacher, argues the answer is *no*. Orators and lawyers sometimes trick people into believing things; sometimes those things happen to be true. People who have been duped in this way have true beliefs, but they don't have knowledge.

Of course, the fact that orators and lawyers intend to deceive isn't crucial for Plato's point. Think of an individual, Clyde, who believes the story about Groundhog



Day. Clyde thinks that if the groundhog sees its shadow, then spring will come late. Suppose Clyde puts this silly principle to work this year. He receives the news about the behavior of the Official Designated Groundhog, and so he believes that spring will come late. Suppose Clyde turns out to be right about the late spring. If there is no real connection between the groundhog seeing its shadow and the coming of a late spring, then Clyde will have a true belief (that spring will come late), but he won't have knowledge.

So what else is needed, besides true belief, for someone to have knowledge? A natural suggestion is that knowledge requires justification. The problem with Clyde is that he didn't satisfy this further requirement (though perhaps he thought he did). Justification, notice, can't just mean that the subject *thinks* he has a reason.

## JUSTIFICATION

What does it mean to say that an individual is "justified" in believing a proposition? Sometimes we start believing that a proposition is true because we consider an argument that describes the evidence available. Mendel believed that genes exist because of the data he collected from his pea plants. Sherlock Holmes believed that Moriarty was the murderer because of the evidence he found at the crime scene. Should we conclude that people are justified in believing a proposition only when their belief in the proposition was caused by their considering an argument? This isn't always plausible. When I believe that I have a headache, I don't construct an argument in which the evidence is laid out in a set of propositions that constitute my premisses. Still, it would seem that I am justified in believing that I have a headache. This suggests that there is such a thing as *noninferential justification*. Some of the propositions we believe are apprehended more or less "directly;" they are not inferred from other propositions that we believe.

What, then, does "justification" mean when it is used in accounts of what knowledge is? When we talk about someone's action being *morally* justified, we mean that the action does not violate any *moral* duties that the person has. Perhaps "justified belief" should be understood in a similar way. We should think of individuals as having certain duties concerning how their beliefs should be formed. A belief is justified if the process by which it was formed does not violate any duties that the person has. To make sense of this suggestion, we'd have to say what duties we have that govern how we are supposed to form our beliefs. Are we obliged to base our beliefs on the evidence that is available, and only on the evidence? This, I take it, is what I do when I believe that I have a headache. However, this suggestion requires further exploration; some of the issues were touched on in Lecture 10, on Pascal's Wager.

Even though the concept of "justification" requires further attention, it is plausible to think that this is one of the necessary ingredients that defines what knowledge is. Knowing that a proposition is true requires more than just having a true belief. The third requirement is that your belief be "justified."

## THE JTB THEORY

Suppose knowledge requires these three conditions. Is that it? That is, are these conditions not just separately necessary, but jointly sufficient? The theory of knowledge that asserts this I'll call the *JTB Theory*. It says that knowledge is one and the same thing as justified true belief:

- (JTB) For any individual *S* and any proposition *p*, *S* knows that *p* if and only if
- (1) *S* believes that *p*.
  - (2) *p* is true.
  - (3) *S* is justified in believing that *p*.

The JTB Theory states a generalization. It says what knowledge is for *any* person *S* and *any* proposition *p*. For example, let *S* be you and let *p* = "the moon is made of green cheese." The JTB Theory says this: If you know that the moon is made of green cheese, then statements (1)–(3) must be true as well. And if you don't know the moon is made of green cheese, then at least one of statements (1)–(3) must be false. As in the definition of bachelorhood discussed in the preceding box, the expression "if and only if" says that we are being given necessary and sufficient conditions for the defined concept.

## THREE COUNTEREXAMPLES TO THE JTB THEORY

In 1963 the U.S. philosopher Edmund Gettier published a pair of counterexamples to the JTB Theory ("Is Justified True Belief Knowledge?" *Analysis*, 1963, Vol. 23, pp. 121–123). What is a *counterexample*? It is an example that goes counter to what some general theory says. A counterexample to a generalization shows that the generalization is false. The JTB Theory says that *all* cases of justified true belief are cases of knowledge. Gettier thought his two examples show that an individual can have justified true belief without having knowledge. If Gettier is right, then the three conditions given by the JTB Theory aren't sufficient.

Here is one of Gettier's examples: Smith works in an office. He knows that someone will soon be promoted. The boss, who is very reliable, tells Smith that Jones is going to get the promotion. Smith has just counted the coins in Jones's pocket, finding there to be 10 coins there. Smith therefore has excellent evidence for the following proposition:

- (a) Jones will get the promotion and Jones has 10 coins in his pocket.

Smith then deduces from this statement the following:

- (b) The man who will get the promotion has 10 coins in his pocket.

Now suppose that, unknown to Smith, Jones will *not* get the promotion. Rather it is Smith himself who will be promoted. And suppose Smith also happens to have 10 coins in his pocket. Smith believes (b); and (b) is true. Gettier also claims that Smith is justified in believing (b), since Smith deduced it from (a). Although (a) is false, Smith had excellent reason to think that it is true. Gettier concludes that Smith has a justified true belief in (b), but Smith doesn't know that (b) is true.

Gettier's other example exhibits the same pattern. The subject validly deduces a true proposition from a proposition that is very well supported by evidence even though it is, unbeknownst to the subject, false. I now want to describe a kind of counterexample to the JTB Theory in which the subject reasons *nondeductively*.

The British philosopher, logician, and social critic Bertrand Russell (1872–1970) described a very reliable clock that stands in a town square. This morning you walk by it and glance up to find out what the time is. As a result, you come to believe that the time is 9:55. You are justified in believing this, based on your correct assumption that the clock has been very reliable in the past. But suppose that, unbeknownst to you, the clock stopped exactly 24 hours ago. You now have a justified true belief that it is 9:55, but you don't know that this is the correct time.

Let me add a third example to Gettier's Smith/Jones story and to Russell's clock. You buy a ticket in a fair lottery. "Fair" means that one ticket will win and every ticket has the same chance. There are 1,000 tickets and you get ticket number 452. You look at this ticket, think for a moment, and then believe the following proposition: Ticket number 452 will not win. Suppose that when the drawing occurs a week later, you are right. Your belief was true. In addition, it was extremely well justified; after all, its probability was extremely close to unity—there was only one chance in a thousand you would be mistaken. Yet, I think we want to say in this case that you didn't *know* that the ticket would fail to win. Here is a third case of justified true belief that isn't knowledge. Note that the reasoning here is *nondeductive*.

## WHAT THE COUNTEREXAMPLES HAVE IN COMMON

In all three of these cases, the subject has *highly reliable*, but not *infallible*, evidence for the proposition believed. The boss *usually* is right about who will be promoted; the clock *usually* is right as to what the time is; and it *usually* is true that a ticket drawn at random in a fair lottery doesn't win. But, of course, *usually* doesn't mean *always*. The sources of information that the subjects exploited in these three examples are *highly* reliable, but not *perfectly* reliable. All the sources of information were prone to error to at least some degree.

Do these examples really refute the JTB theory? That depends on how we understand the idea of justification. If highly reliable evidence is enough to justify a belief, then the counterexamples do refute the JTB theory. But if justification requires perfectly infallible evidence, then these examples don't undermine the JTB theory.

My view is that justifying evidence needn't be infallible. I think we can have rational and well-supported beliefs even when we aren't entitled to be absolutely certain that

what we believe is true. From this, I conclude that justified true belief isn't sufficient for knowledge.

## AN ARGUMENT FOR SKEPTICISM

What more is required? The lottery example suggests the following idea. In this case, you *probably* won't be mistaken when you believe that ticket number 452 will lose. But high probability isn't enough. To know that the ticket will lose, it must be *impossible* for you to be mistaken. You don't have knowledge in this example, because there is a chance (small though it may be) that you will be wrong.

There is some plausibility to this suggestion about what knowledge requires. The problem is that this idea seems to lead immediately to skepticism—to the conclusion that we don't know anything. For it seems that virtually all the beliefs we have are based on evidence that isn't infallible. Consider, just briefly, the beliefs we have that depend on the testimony of our senses. We use vision, hearing, touch, and so forth, to gather evidence about the way the world is. Do the resulting beliefs count as knowledge? The problem is that the senses are sometimes misleading.

Right now, you believe you are looking at a printed page. You believe this because of the visual experiences you now are having. Do you *know* that there is a printed page in front of you? According to the present suggestion, for this to be true, it must be the case that you couldn't possibly be mistaken in believing what you do. But the fact of the matter seems to be that you *could* be mistaken. You might be hallucinating, or dreaming, or your senses might be malfunctioning in some other way.

So here's where we are. The JTB Theory is mistaken. JTB doesn't suffice for knowledge. As an alternative to the JTB theory, we have this suggestion: Knowledge requires the impossibility of error. But this suggestion, plausible though it may be as a diagnosis of why you don't have knowledge in the lottery example, allows us to formulate the following argument for skepticism:

If *S* knows that *p*, then it isn't possible that *S* is mistaken in believing that *p*.

It is possible that *S* is mistaken in believing that *p*.

---

*S* doesn't know that *p*.

This form of argument can be used to argue that a posteriori knowledge is impossible. Recall from Lecture 8 that a proposition is a posteriori if it can be known only through the testimony of sense experience. The skeptic claims that beliefs based on sense experience aren't totally immune from the possibility of error. People make perceptual mistakes—for example, in cases of illusion, hallucination, and dreaming. Since we can't absolutely rule out the possibility of error, the skeptic concludes that we must admit that the senses don't provide us with knowledge.

### Is Skepticism Self-Refuting?

Skeptics claim that people don't know anything. Can skeptics claim to know that what they say is true? If not, does that show that their philosophical thesis is false?

It is a contradiction to say that you know that no one (including yourself) knows anything. But a skeptic can assert that people lack knowledge without claiming to know that this is so. In addition, a skeptic can claim to provide a good argument for skepticism. This isn't contradictory.

The thesis that no one knows anything, if true, can't be known to be true. But that doesn't show that the position is false. If you think that all truths are knowable, you will say that this skeptical thesis can't be true. But why think that all truths are knowable? Why not think, instead, that the universe may contain truths that we are incapable of knowing?

Another way for skeptics to avoid contradicting themselves is to be modest. Instead of claiming that no one knows *anything*, they could limit themselves to the claim that no one ever knows anything *through the testimony of the senses*. If this more limited kind of skepticism could be supported by an a priori argument, the position wouldn't be self-refuting.

This argument for skepticism is deductively valid. (Identify its logical form.) The first premiss seems to describe a plausible requirement for knowing the proposition in question. The second premiss also seems plausible; it just says that the beliefs we have aren't absolutely immune from the possibility of error. If you want to reject skepticism, you must refute one or both of these premisses.

Although the above argument has fairly plausible premisses, its conclusion is pretty outrageous. I think I know lots of things, and I believe this is true of you as well. It is hard for me to accept the idea that I don't know that I have a hand. If you think that you now know there is a printed page in front of you, you should balk at this argument as well. The skeptical argument contradicts a fundamental part of our commonsense picture of the way we are related to the world around us. Common sense says that people have knowledge of the world they inhabit; the skeptical argument says that common sense is mistaken in this respect. If there is a mistake in the argument, where is it?

---

### Review Questions

1. Explain how object knowledge, know-how knowledge, and propositional knowledge differ.
2. What do the following pieces of terminology mean? (i) *X* is a necessary condition for *Y*; (ii) *X* is a sufficient condition for *Y*; (iii) *X* is true if and only if *Y* is true; (iv) *o* is a counterexample to the statement "All emeralds are green."
3. Why think that true belief isn't sufficient for knowledge? What is the JTB theory? What is the difference between highly reliable evidence and absolutely infallible evidence?

4. Describe the three counterexamples to the JTB Theory (Gettier's, Russell's, and the lottery example). Do these counterexamples show that JTB isn't necessary or that it isn't sufficient for knowledge?
5. What is skepticism? What is the argument for skepticism given at the end of this lecture? Is the argument deductively valid?

---

### Problems for Further Thought

1. People sometimes say, about a proposition that is difficult to take to heart, "I know it, but I don't believe it." Does this refute the claim that if  $S$  knows that  $p$ , then  $S$  believes that  $p$ ?
2. I suggested that we lack knowledge in the lottery example because knowledge requires the impossibility of error. Can you think of some other explanation for why we lack knowledge in this or in the other two examples that were used against the JTB Theory?

## LECTURE 13

---

### *Descartes's Foundationalism*

René Descartes (1596–1650) is sometimes described as the father of modern philosophy. The kind of epistemology he tried to develop is called foundationalism. Before launching into the details of Descartes's philosophy, I want to describe what kind of approach to the problem of knowledge foundationalism provides.

### FOUNDATIONALISM

The word *foundationalism* should make you think of a building. What keeps a building from falling over? The answer has two parts. First, there is a solid foundation. Second, the rest of the building, which I'll call the superstructure, is attached securely to that solid foundation. Descartes wanted to show that (many if not all of) the beliefs we have about the world are cases of genuine knowledge. To show this, he wanted to divide our beliefs into two categories. There are the foundational beliefs, which are perfectly solid. Then, there are the superstructural beliefs, which count as knowledge because they rest securely on that solid foundation.

Besides the metaphor from architecture, there is another that should help you understand what Descartes's project is. You probably had a geometry course in high school. Here you studied Euclid's development of the subject. Recall that Euclid, who lived about 2,200 years ago, divided the propositions of geometry into two categories.