THE GREAT COURSES®

Topic
Philosophy & Intellectual History

Subtopic
Applied Philosophy

Theories of Knowledge
How to Think about What You Know

Course Guidebook

Professor Joseph H. Shieber
Lafayette College

Published by
THE GREAT COURSES

Corporate Headquarters

4840 Westfields Boulevard | Suite 500 | Chantilly, Virginia | 20151-2299

[PHONE] 1.800.832.2412 | [FAX] 703.378.3819 | [WEB] www.thegreatcourses.com

Joseph H. Shieber, PhD

Associate Professor of Philosophy
Lafayette College

Joseph H. Shieber is an Associate Professor in the Department of Philosophy at Lafayette College, where he has taught since 2003. Before arriving at Lafayette, he taught philosophy at Brown University, Connecticut College, and the Massachusetts Institute of Technology. Dr. Shieber earned a BA in Literature from Yale University, studied mathematics and philosophy at the Freie Universität Berlin, and earned AM and PhD degrees in Philosophy from Brown University.

Dr. Shieber has published numerous articles in epistemology, philosophy of language, and the history of modern philosophy in some of the top journals in the field of philosophy. He is also the author of *Testimony: A Philosophical Introduction* as well as a monthly Monday columnist for *3 Quarks Daily* (www.3quarksdaily.com).

Dr. Shieber is a recipient of the Thomas Roy and Lura Forrest Jones Faculty Lecture Award in recognition of excellence in teaching and scholarship from Lafayette College. He regularly teaches courses in theory of knowledge, philosophy of language, history of 20th-century philosophy, logic, and metaphysics as well as a first-year seminar on propaganda.

Table of Contents

INTRODUCTION

GUIDES

SUPPLEMENTARY MATERIAL

NAVIGATION TIP

To go back to the page you came from, **PRESS ALT** + ← on a PC or ⌘ + ← on a Mac. On a tablet, use the bookmarks panel.

THEORIES OF KNOWLEDGE

How to Think about What You Know

These lectures are intended as an introduction to the philosophical analysis of knowledge. You can think of them as organized into roughly four sections.

1. The section on the basic theories of the structure of knowledge—the theories lectures—comprises lectures 1 through 6.
2. The section on the personal sources of knowledge—the personal sources lectures—spans lectures 6 through 16.
3. The section on the social source of knowledge—the social sources lectures—encompasses lectures 17 through 21.
4. The section that introduces some special topics and cutting-edge research in the study of knowledge—the special topics lectures—runs from lectures 22 to 24.

The theories lectures discuss the main divisions in philosophical theories of knowledge. Roughly, there are two divisions that are explored.

- The first division is between theories that seek to explain knowledge as a hierarchical system of beliefs, with more-certain beliefs serving as a foundation for other, less-certain beliefs; and theories that seek to explain knowledge as an interconnected, mutually supporting web of beliefs. The first type of theories are known as foundationalist theories while the second type are called coherentist theories.
- The second division is between internalist theories and externalist theories. Internalist theories stipulate that the resources that knowers can appeal to in order to support their knowledge consist solely of mental states, states that are in principle available to the awareness of the knowers themselves. In contrast, externalist theories allow that at least some of the resources on which knowers depend in order to support their knowledge involve states of the brain, body, or perhaps even the world and that knowers don't need to have any access to those states in awareness.

After getting that theoretical heavy lifting out of the way, the course turns to a discussion of different sources of knowledge. The personal sources lectures discuss different sources of knowledge that have traditionally been thought to depend only on the resources of individual knowers.

These include your knowledge of your own internal mental states, your knowledge on the basis of sense perception, your memory knowledge, your knowledge on the basis of reason, and the knowledge that you derive through inference. This group of lectures concludes with a discussion of the topic of know-how, or ability knowledge.

The social sources lectures tackle the topic of knowledge acquired through communication with others. The philosophical term for such knowledge is testimony. Though the topic of testimony was ignored for much of the history of Western philosophical thought, it has recently enjoyed a resurgence of interest. The lectures on testimony examine three theories of testimony that are currently popular. You will discover that there is a great deal of data from social psychology that casts doubt on whether any of those theories can be workable, and you will be introduced to a theory that better fits the data. Then, that theory will be illustrated by applying it to discussions of how the sciences and the media contribute to our acquisition of knowledge.

The special topics lectures consider a few areas of cutting-edge interest in the theory of knowledge. These include the topic of pragmatic and moral encroachment—roughly, whether our pragmatic or moral commitments could influence whether we acquire knowledge on the basis of the evidence that we have. The lectures also include the age-old topic of skepticism and a recent philosophical theory, contextualism, that was formulated to deal with the skeptical challenge. The lectures end with a consideration of the future of philosophical reasoning about knowledge.

There is a wealth of rich topics and puzzles concerning the nature of knowledge, any number of which could serve as the theme for a whole series of lectures in their own right. At the conclusion of these lectures, you'll be well positioned to explore further to discover more of the fascinating features of any of the topics in the discussion of knowledge that interest you.

LECTURE 1

Philosophy and Transformative Experiences

Contrary to what you might think, the theory of knowledge couldn't be more topical nowadays. Within the last week, have you heard a report, read an article, or had a conversation about fake news? What about artificial intelligence? Or big data? All of these pressing issues in contemporary society have to deal with knowledge: how we acquire it, how we preserve it, how we communicate it. And thinking philosophically about knowledge will help you formulate rational solutions to real-world problems.

RATIONAL DECISION MAKING

- A recent challenge to common views about the extent of our knowledge and rationality was formulated by philosopher L. A. Paul in her book *Transformative Experience* and an accompanying article, "What You Can't Expect When You're Expecting."

- Paul asks us to consider a common way of thinking about transformative decisions—for example, the decision whether or not to have a child. It is a common piece of advice to suggest that you should approach this weighty decision carefully. In particular, people usually recommend that you decide in large part on the basis of considering what it would be like to have a child.

- In fact, this seems to be a paradigm of rational decision making. First, you need to ensure that you are in a situation in terms of your maturity, financial security, and physical health to have a child. After that, you and your partner imagine what it would be like to have a child, as well as what it would be like to remain childless, and compare the two possible futures. Only then—when you have decided which of those possible scenarios seems to promise a more fulfilling and happy life—would you choose which course of action to pursue.

- There seem to be two principles here:
  - Transformative experience: Some experiences (such as becoming a parent for the first time) change your life in a fundamental way.
  - Expected experiential value: In deciding whether to pursue transformative experiences, you should make that decision based on your expectation of the overall value of such experiences.

- Suppose there really are transformative experiences. The problem is that, if there are, then making decisions about such experiences based on what you expect the value of such experiences to be is irrational. In other words, the two principles that we formulated are mutually inconsistent.

- There are really two separate problems with using your imagination about future experiences to decide whether to pursue transformative experiences.
  - The first problem is perhaps easier to explain, but it is actually the less significant challenge. If an experience is genuinely novel, then you won't in fact be able to use your imagination to know what that experience is like. In the case of these future experiences that you have absolutely no knowledge about, it's hard to see

Anecdotally, there has long been a recognition that the love of a parent for his or her child changes the parent in profound ways. But we don't need to rely only on anecdotal evidence of the transformative effect of becoming a parent. Scientists have been able to trace such effects in the physiology of the brain.

In a 2015 review article in the journal *Cell*, Dr. Ruth Feldman discusses the ways that the brains of both men and women change as a result of becoming new parents. She notes that such changes include not only hormonal changes, but also changes in the activity of the brain, as measured by fMRI. Dr. Feldman takes such brain activity to be evidence of a "parental caregiving" network in the brain.

how imagination can give you any knowledge of what such experiences are like.

- The second problem is not as easy to state, but it's actually the more serious difficulty for the experiential knowledge view. The problem is that transformative experiences change the way that you value the experiences that you acquire.

- What it means for an experience to be transformative is in part that it changes the way you conceive of a good life. In other words, having that transformative experience will change what things you take to be part of a good life and how you would rank those things in order of importance. For this reason, it's not possible to decide whether to pursue a life-changing experience by imagining how that experience will bear on the quality of your life.

- Philosopher L. A. Paul has convincingly argued that two commonsense principles for making rational decisions are in conflict. It can't be the case both that some experiences are genuinely transformative and that rationality demands that we decide whether to pursue transformative experiences by imagining what those experiences are like. Does that mean that there is no rational way to decide whether to pursue such experiences?

- Let's consider the full implications of Paul's challenge. Initially, you might wonder how terrible it would be if choices about transformative experiences were limited to a few dramatic cases like the choice whether to become a parent. However, if the conclusion we've arrived at so far is all we're left with, that would in fact be a terrible result. This is because even though the choice as to whether to become a new parent is a particularly dramatic example of a transformative experience, there are many other experiences that must also be described as transformative—such as whether to move to a new city or whether to date or marry a certain person.

- Economists speak of opportunity costs: the costs that result from the fact that pursuing one course of action means that you can't simultaneously pursue other, incompatible courses of action. Often, when those opportunity costs involve the threat of missing out on new experiences, your decisions about acting will result in the kind of troubling situation that Paul describes. In other words, our lives are filled with many such situations in which we must decide whether to pursue what Paul calls transformative experiences, so the threat of not being able to decide rationally in such situations is a very great threat.

THE IMPERSONAL STRATEGY

- Trying to decide whether to pursue a transformative experience by imagining that experience and then deciding whether that imagined experience is, all things considered, worth pursuing is not rational. There is, however, another way that we might try to decide whether to pursue a transformative experience.

- Rather than imagining the experience, we could instead try to gather evidence about people who have pursued the transformative experience. Then, we could see whether those who have pursued the experience are better off than those who have not pursued the experience. In other words, we could completely ignore our subjective experiences and preferences and try to gather evidence impersonally in order to arrive at a decision.

- It turns out that, in the case of deciding whether to become a parent, we might not yet have good enough evidence to allow us to reach a decision in this way. When comparing the happiness of two different groups, you have to control for other features that contribute to happiness, such as income, health, and having a committed romantic relationship. Once you control for those factors, most studies suggest that there is no difference in happiness between parents and childless couples. Indeed, some studies suggest that, if anything, childless couples—in the United States, at least—are happier than couples with children living at home.

- Let's leave aside these issues, however. There's a deeper worry for those who might encourage us to adopt the impersonal way of deciding about the pursuit of transformative experiences. Consider the following case:
  1. Suppose the evidence is clear that couples who have children are happier than those who don't.
  2. And suppose you are a person who, when you imagine what it would be like to have children, finds no appeal in that imagined experience.
  3. Suppose furthermore, however, that you're familiar with Paul's argument about the irrationality of using imagined experience to decide whether to pursue transformative experiences.
  4. Finally, suppose you're very committed to employing a rational procedure to decide whether to become a parent.

- In such a case, it would seem that the rational course of action would be to decide to have a child. And this is true despite the fact that the prospect of becoming a parent holds no appeal for you.

- The good news is that there is, in fact, a way of deciding available to us—the impersonal strategy—that is rational. The bad news is that now, according to Paul, we're faced with a dilemma: two options, neither of which is appealing. We have a commonsensical way of deciding that has turned out not to be rational and a rational way of deciding that conflicts with common sense!

- Paul is certainly correct that many people would find it funny to decide based on impersonal decision making. Most people would rather rely on their own subjective impressions when making important life decisions. It can be argued, however, that our goal should be to make people more comfortable with employing decision-making strategies that are, in fact, rational, rather than emphasizing the counterintuitive aspects of such rational decision-making strategies.

- Indeed, there has now been a great deal of evidence from social psychology regarding the ability of human beings to predict their future happiness on the basis of their imagined future experiences. Psychologists call such predictions affective forecasting. Suppose we forget about whether it's impossible to decide rationally about transformative experiences on the basis of imagining our future experiences. Social psychologists have gathered a great deal of evidence that the fact is that we are very poor affective forecasters when we rely on our imagined future experiences.

In a 2009 paper in the *Journal of Experimental Psychology: Applied*, psychologists Emma Walsh and Peter Ayton present a series of experiments to demonstrate that, when given access to both forms of evidence—impersonal reporting and subjective imagining—those who ignored their own imagined experiences performed better at predicting their future experiences.

In fact, studies have shown that if you're given both types of evidence—evidence that would allow you to imagine your future experience as well as evidence about the reports of strangers who have had the experience already—you would do better deciding simply on the basis of the strangers' reports. This suggests that our imaginations about our own experiences are simply sources of error that we should avoid, rather than evidence that ought to be taken into account—even in combination with impersonal reporting.

Three lessons from this discussion can be drawn that will carry forward in subsequent lectures.

- Common sense is not always an accurate guide as to which strategies are actually useful in arriving at accurate information or reaching rational decisions.
- Subjective experience is not necessarily a source of useful evidence upon which to base your beliefs or decisions.
- Abstract, philosophical discussions can be enriched by considering evidence from empirical disciplines—for example, in the case of transformative experiences, from social psychology.

READINGS

Musgrave, *Common Sense, Science and Scepticism*.

Paul, *Transformative Experience*.

QUIZ

1 TRUE OR FALSE

The most rational way to decide whether to pursue some experience is to imagine that experience and then to see whether what you imagined is something that you'd like to pursue.

2 Which of these are challenges to using expected experiential value to make decisions about whether to pursue transformative experiences?

- **a** If an experience is genuinely transformative, you can't imagine in advance what that experience is like.
- **b** A transformative experience also changes how you value your experiences, so the significance that you assign to the experience itself can change.
- **c** Both a and b.
- **d** Neither a nor b.

3 TRUE OR FALSE

Even if it's not possible to decide about transformative experiences using expected experiential value, at least transformative experiences are not very common.

TIP: To go back to the page you came from, **PRESS ALT** + ← on a PC or ⌘ + ← on a Mac. On a tablet, use the bookmarks panel.

4 According to experimental results in social psychology, which method is the best way for you to predict your future happiness?

- **a** Simulating a future experience and imagining how it would make you feel
- **b** Relying on other people's reports about their experiences
- **c** Weighing your own simulation of the future experience as well as others' reports about their experiences and coming to some judgment that combines the two

5 According to Plato's allegory of the cave, why is it difficult for philosophers to convey their knowledge to others?

- **a** Philosophers have trouble adjusting their eyes after transitioning from the bright light of truth to the dim shadows of the cave.
- **b** Those who remain in the cave have better knowledge of the shadows cast by the puppets on the cave walls.
- **c** Both a and b.
- **d** Neither a nor b.

CLICK to navigate.

Answer key can be found on page 207.

LECTURE 2

Knowledge, Truth, and Belief

Philosophers have been studying knowledge for thousands of years. In fact, one of the writings on knowledge that is still read today was written around 369 BC. It's one of Plato's Socratic dialogues, the *Theaetetus*, in which Plato has Socrates discuss the nature of knowledge with some young Athenians, among them the nobleman Theaetetus. At the end of the dialogue, Plato, speaking through the character Socrates, arrives at a definition of knowledge as "true belief with *logos*," the Greek word meaning "a reason" or "an account."

THE CONNECTION BETWEEN KNOWLEDGE AND TRUE BELIEF

- There is a great deal of evidence from developmental psychology that the ability to recognize instances of beliefs and knowledge in others is a fundamental ability for children. In fact, this ability is essential to what psychologists call a child's theory of mind—their ability to use reasoning about other people's mental states to explain and predict their behavior.

- Research into the development of theory of mind in children has exploded since the 1980s, in no small part due to the initial work of Austrian developmental psychologists Josef Perner and Heinz Wimmer. Their first experiment, published in *Cognition* in 1983, has now been reproduced numerous times, with strikingly consistent results. The widely accepted implication of these results is that children at the age of three to four years are unable to form beliefs concerning the false beliefs of others.

- One reason why the results regarding false belief are particularly significant is because of a connection they suggest between understanding belief and understanding truth and knowledge. It turns out that children's abilities to make sense of belief develops at about the same time as their abilities to understand what it is to deceive someone intentionally. This would suggest that children begin to acquire a more sophisticated understanding of the meaning of truth at this stage.

- It is also at this same time that children develop the ability to recognize the connection between a person's sources of information and that person's knowledge.

- In his 2012 book *Trusting What You're Told: How Children Learn from Others*, the Harvard developmental psychologist Paul Harris spells out a number of ways in which children between the ages of three and six begin to be sensitive to which informants are knowledgeable.

- As Harris argues, even newborn children look to adults to gain information both about their environment and how to act. At the earliest stages, however, children base their trust on social and emotional cues: Am I familiar with this adult? Do I already have a positive emotional attachment to him or her?

- Beginning in the preschool years, however—in that crucial transition from ages three to six that Wimmer and Perner also explore—children begin to search more for evidence of accuracy and reliability when deciding who to trust: Was the person in a good position to perceive the event he or she is reporting about? Has the person been accurate in the past? Does the person use names for objects that I'm familiar with in the way that I expect they should?

- Wimmer and Perner offer a good summary of these connections:

> *[A] novel cognitive skill seems to emerge within the period of 4 to 6 years. Children acquire the ability to represent wrong beliefs and to construct a deceitful or truthful utterance relative to a person's wrong beliefs. Within this period several other related abilities also emerge: children start to understand another person's absence of knowledge … . They become able to construct a deceitful utterance … and to infer a deceptive plan from a critical utterance in the context of conflicting goals.*

Plato's 2,000-year-old theorizing and contemporary developmental psychological research agree: There's an intimate connection between knowledge and true belief.

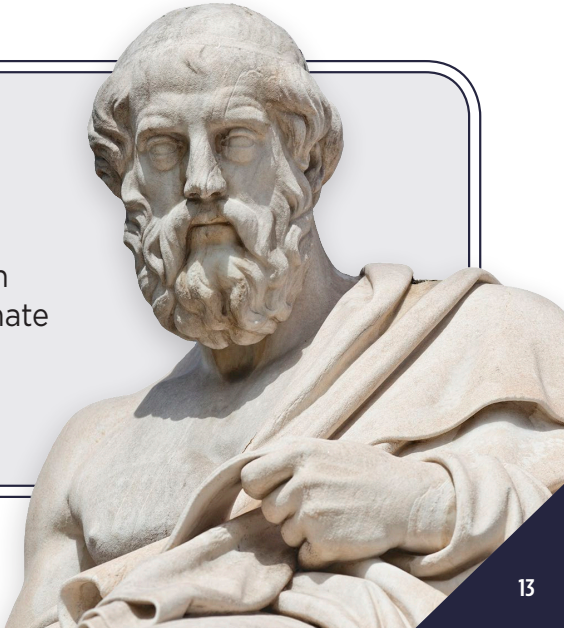

INEXISTENCE AND INACCESSIBILITY ARGUMENTS

- The analysis of knowledge as involving true belief faces a significant challenge, however. There's a very commonly held notion that belief—being by its nature a subjective, individual representation of reality—is incompatible with objective truth.

- In fact, contemporary society seems to be full of controversies that rest on different perspectives on the fundamental objectivity of truth. In a 2017 essay in *The Chronicle of Higher Education* entitled "Teaching Humility in an Age of Arrogance," philosophy professor Michael Lynch criticizes the idea that there is no such thing as truth. He puts the argument against the notion of objective truth this way:

> *There just is no way of escaping your perspective or biases. Every time you try to get outside of your own perspective, you just get more information filtered through your own perspective. As a consequence, objective truth is just irrelevant—either we'll never know it or it doesn't exist in the first place.*

- Unfortunately, although Lynch does not endorse this argument against the relevance of objective truth, he does not in his essay explain what is wrong with the argument. To be fair to Lynch, he's more concerned with arguing in favor of intellectual humility and an openness to evidence.

- However, it is a shame that he doesn't debunk the argument against objective truth. It's one that seems to be very appealing to many people, perhaps because they associate the notion of objective truth with dogmatism or close-mindedness. That association, however, is a mistake. In fact, the idea that it's possible that we could be very confident and believe ourselves to have very good evidence but nevertheless be mistaken about the truth is the opposite of dogmatism. But that possibility is exactly what objective truth guarantees! So, arguing for the notion of objective truth serves Lynch's broader goal of making the case for intellectual humility.

- The quote from Lynch contains two arguments:
  - One argument begins with the idea that we always only have access to information from our own subjective impressions and concludes with the claim that objective truth doesn't exist. Let's call this the inexistence argument.

- The other argument also begins with the idea that we always only have access to information from our own subjective impressions but concludes with the claim that we'll never know objective truth. Let's modify the conclusion of this argument slightly by making it that we'll never *grasp* objective truth. This modification doesn't really affect the point of the argument; it doesn't seem to be an argument about knowledge. It's really about truth and its inaccessibility, so let's call this the inaccessibility argument.

- The problem with the inexistence argument is that it involves a pretty obvious non sequitur, which is a fancy way of saying that the conclusion of the argument doesn't follow from its premises. In this case, the glaring problem is that how we access information has no bearing on whether or not truth exists.

- Suppose someone argues that, because architectural photography always only records information from the subjective impressions of the photographer, buildings don't exist. This is ridiculous. But it's exactly what the inexistence argument attempts to do in the case of truth.

- So, the inexistence argument fails: There is no reason to think that objective truth does not exist, even if—for the sake of argument—we grant that we always only have access to our own subjective impressions.

- We'll allow, then, for at least the possibility that truth exists. For our purposes, we may characterize truth as involving a certain kind of state of affairs—for example, a state of affairs consisting of a claim and something the claim is about. And, to continue the example, what it is for a claim to be true is for what is claimed actually to be the case. So, the claim "Washington wintered at Valley Forge in 1777–1778" is true so long as Washington really did winter at Valley Forge in 1777–1778.

- What is a claim? It's something that some individual person makes. There can't be a claim without someone claiming it—even if that person is only claiming it, say, for the sake of argument. States of affairs, however, are not claims. They are ways that the world is. And, as such, they are independent of claims.

- States of affairs, then, are objective and independent of anybody's thoughts or claims. Claims, on the other hand, are subjective; they are acts—speech acts—that an individual person performs. So, what it means for something to be true *always* involves a subjective component. And what it means to grasp an objective truth is simply to endorse a claim that is true.

There are at least three distinct kinds of knowledge.

- Know-how is the knowledge that craftspeople have; they know how to pilot a ship or cure disease.
- Knowledge-wh—knowledge-who, -what, -where, or -why—is the knowledge that enables you to answer a who, what, when, where, or why question.
- Knowledge-that, or propositional knowledge, is knowledge that a certain fact is true.

Most philosophical discussions of knowledge have traditionally focused on propositional, or factual, knowledge—knowledge-that.

- For our purposes right now, we can treat beliefs as on a par with claims; roughly, beliefs are claims that you endorse silently to yourself, or claims that you would endorse silently to yourself if they were brought to your attention, or claims that you would endorse publicly if they became a topic of interest.

- So, if our argument shows that the subjectivity of claims is completely compatible with our ability to grasp objective truth, then the argument should equally show that the subjectivity of beliefs is completely compatible with our ability to grasp objective truth.

- In other words, rather than its being incompatible with the subjectivity of belief that we grasp objective truth, our grasping objective truth always involves an aspect of subjectivity.

ARGUMENT FOR THE IRRELEVANCE OF TRUTH

- This argument has been proposed in slightly different versions by a number of thinkers, but the version we'll consider is that offered by philosopher Richard Rorty, who begins by noting that, whenever we ask ourselves whether some claim is true, we answer that question by looking at the evidence that we have for thinking that it is true. If that's the case, he continues, then there is no way to determine truth independently of our evidence. So, he suggests, truth is irrelevant—all that we really ought to concern ourselves with is our evidence for endorsing a given claim.

- There are reasons to question the premises of Rorty's argument, but for now we don't even need to do that. This is because the biggest problem with Rorty's argument is that it is a non sequitur. As with the inexistence argument, in Rorty's argument the conclusion simply doesn't follow from the premises.

- We can see this most clearly by considering how Rorty is thinking of truth for his argument to work. In order for the argument to work, the only use of truth can be when we ask *ourselves* whether something is true. But this is not the only use. Sometimes you want to know not whether you think something is true, but whether someone else thinks something is true.

- Once you consider those sorts of cases, it's obvious that grasping the truth and having evidence for the truth can come apart—in either direction. In other words, sometimes you can tell that someone happened upon the truth but believes it on the basis of not very good evidence; other times you can see that someone thinks he or she has good evidence that something is true, but you can tell that it isn't true. So, Rorty's argument for the irrelevance of truth fails as well.

READINGS

Alston, *A Realist Conception of Truth*.

Audi, *Epistemology*.

QUIZ

1 Which of the types of knowledge involves factual knowledge?

a Know-how

b Knowledge-wh

c Propositional knowledge

2 What definition of knowledge does Socrates arrive at by the conclusion of the *Theaetetus*?

a True belief with techne

b True belief

c Truth

d True belief with logos

3 TRUE OR FALSE

According to Josef Perner and Heinz Wimmer, most three- to four-year-olds already possess a well-developed theory of mind.

4 TRUE OR FALSE

The subjective component of belief ensures that there is no possibility for accessing objective truth.

5 TRUE OR FALSE

Heinz Wimmer and Josef Perner's false belief task provides us with evidence against Richard Rorty's argument that truth is irrelevant.

Answer key can be found on page 207.

LECTURE 3

Foundationalism: Descartes's Evil Demon

In 1619, René Descartes experienced a revelation that caused him to seek to find a way to discover a "universal method for the discovery of truth." The method that he propounded is considered by many scholars—and many of the thinkers who came after Descartes—to be the beginning of modern philosophy. Descartes compares the sciences as they have traditionally been practiced to a building erected on a shaky foundation. Just as we wouldn't continue to renovate and improve a building that we knew to have a shaky foundation, Descartes suggests that we shouldn't continue to increase the store of our scientific knowledge until we have shored up its foundation.

THE EVIL DEMON ARGUMENT

- One way to think about the "universal method" that Descartes discovered is the way that Descartes himself later argued for the implementation of his method, which has come to be known as the evil demon argument.

- Descartes begins by noting that often, in dreams, we experience things that aren't real. Within dreams, however, we believe that the things we are dreaming are real. Of course, even in dreams we might still be said to know very general facts about the world—for example, that the earth revolves around the sun or that fire is hot. If his goal is to find a way to find a new basis for all of the sciences, then, Descartes knew that he would need a way to tear down the entire edifice of knowledge so that he could rebuild it from the ground up using his new method.

- To do this, Descartes asks us to imagine a very powerful demon, one so powerful that he can not only make us experience things that aren't there, as we do in particularly vivid dreams, but also to believe things that aren't true—even that 2 + 2 = 5, for example.

- If we imagine that we are in the clutches of such a powerful demon,

- This is the origin of Descartes's well-known assertion *Cogito ergo sum*: "I think, therefore I am." Even if the demon can deceive you about the content of your beliefs, he cannot deceive you about the fact that he is now causing you to doubt those beliefs. If you are doubting, then you are thinking. If you are thinking, then you exist. Using Descartes's new method, we have now established definitively the fact of our own existence upon the firmest of all possible foundations.

INTERNALISM AND DUALISM

- As the great logician, philosopher, peace activist, and Nobel Prize winner Bertrand Russell writes regarding Descartes,

> *There is thus, in all philosophy derived from Descartes, a tendency to subjectivism, and to regarding matter as something only knowable, if at all, by inference from what is known of the mind.*

- As this quote from Russell's *A History of Western Philosophy* illustrates, Descartes's theory of knowledge is characterized by two components: what Russell terms "subjectivism" regarding the source of the beliefs that serve as the foundations of knowledge and how all of the other knowledge that we have derives from the foundations by means of inference. This lecture will focus on the first component, while the next lecture will focus on the second.

- Philosophers since Russell would quibble with the use of the term "subjectivism" to characterize Descartes's position. Instead, contemporary usage would characterize Descartes as an internalist, because according to Descartes, all of the evidence upon which we build our edifice of knowledge is internal to the mind of the knower.

- At the core of Descartes's method is a focus on what you are experiencing right now, in the moment. In other words, Descartes's method builds the foundation for all knowledge out of the states of your mind at the given instant in which you are experiencing those states—what philosophers call occurrent mental states, because the states are only experienced as they are occurring. So, what Descartes is really saying is that all of our knowledge ultimately rests on our occurrent mental states for its foundation and that these occurrent mental states are more certain for us

than any other information about which we might form beliefs.

- It is important to stress that the sort of evidence available to the internalist like Descartes is what is present to the mind—not what occurs in the brain.

- As Descartes notes, a mind is a thinking thing, one that "doubts, understands, conceives, affirms, denies, wills, refuses, which also imagines and feels." What the mind feels, however, is distinct from any information that, say, the human sense organs gather.

- This is obviously true for Descartes himself. Descartes was what philosophers call a dualist, which means that he believed that minds and brains are two separate and distinct substances. In particular, Descartes believed that minds had no spatial location—though they did interact with the brain.

- Few philosophers are Descartes-style dualists anymore. Many philosophers, however, still subscribe to internalism about the ultimate sources of knowledge.

Strangely, Descartes thought that mind-brain interaction took place at the pineal gland, a small gland at the center of the brain, right between the two halves of the thalamus.

Scientists now know that the pineal gland produces melatonin, a hormone essential to maintain our sleep cycle, although in Descartes's day, the true function of the pineal gland was unknown.

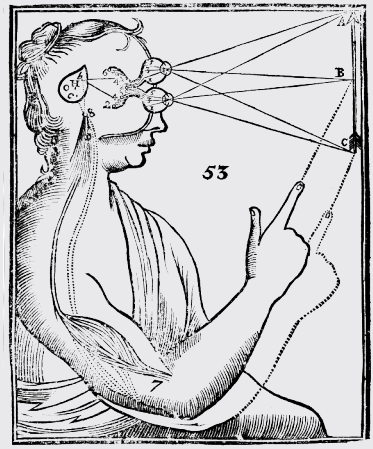

- Let's suppose that contemporary internalists agree with Descartes that the foundational evidence on which all knowledge rests concerns occurrent mental states and that they disagree with Descartes about the separation of mind and body. So, let's say that they either think that the mind just is some part of the body—some part of the brain, perhaps—or that the mind is the result of the operation of some part or parts of the body.

- Clearly, what constitutes the mind doesn't necessarily involve *all* of the body. A limited number of parts of the brain make up the mind.

- Contemporary theories of knowledge that embrace internalism—even those that don't follow Descartes in radically separating mind from body—still must limit the sources of evidence. Those sources that serve as the foundations of knowledge for the internalist will be limited to occurrent mental states that are at best a very small subset of the larger number of states of the brain, central nervous system, or peripheral nervous system.

- Descartes's view of knowledge requires that our occurrent mental states be the most certain type of knowledge that we can have. This is something we can test: Is it plausible that our knowledge of our occurrent mental states reaches the level of certainty that Descartes's theory requires?

- There are a number of different types of occurrent mental states. Descartes himself speaks of doubting, understanding, conceiving, affirming, denying, willing, refusing, imagining, and feeling. It isn't very plausible that our knowledge of all of these mental states will be equally certain.

- For now, let's give Descartes the benefit of the doubt and focus on the most promising case for certain knowledge of occurrent mental states: feelings—in particular, the feeling of having a certain visual sensory experience.

- Phenomena from vision science—such as foveal vision, saccades, change blindness, and inattentional blindness—demonstrate that if we were forced to rely on our awareness of our occurrent visual experiences, we would be in trouble. These and other phenomena demonstrate that in the case of vision, we have very little awareness of the quality of our occurrent visual experiences.

- If this is true of vision, then there is little hope of Descartes's strategy working for other forms of sensory awareness either. For most people, visual experiences are our most vivid and rich. If, as the evidence shows, we are misled about the extent of our awareness of visual experience, this does not suggest that we should be optimistic about our awareness of the phenomenal qualities of sounds, smells, tastes, or touches.

Reasons for skepticism about the certainty of our awareness of occurrent sensory experiences come from the field of vision study.

You've probably heard of peripheral vision, but you may not have heard of its complement: foveal vision. The fovea is a small region of very tightly packed cones, the cells responsible for color and shape detection in the eye. Almost half of all the information that reaches the brain comes from this one region, and only the part of the scene captured by foveal vision is seen with any detail.

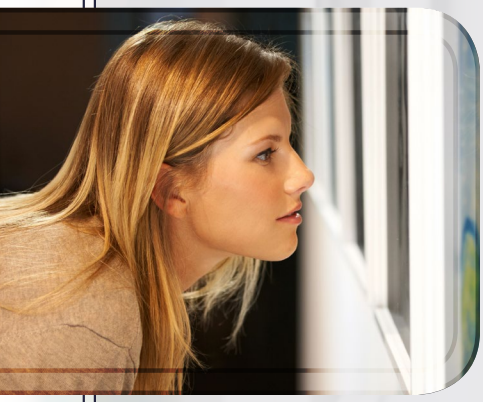

We're unaware of this, of course. When we think about our visual experience, it seems to us as if we see an entire scene in rich detail. This is because our eyes are constantly moving, even when we're examining a still image. These movements, called saccades, are extremely fast. They allow our brain to gather detailed information about different parts of the visual scene as our foveal vision scans it for interesting features, all of which occurs without our conscious awareness or control.

The area captured by foveal vision is extremely small—the width of one thumbnail when viewed at arm's length. So, the detailed visual scenes that we seem to experience are actually composed of a number of these thumbnail-sized images captured by our eyes, which constantly move, or saccade, to take in different areas of the visual scene. Our brains then piece these together to form a larger whole.

In the case of change blindness, there are many features of the visual scene that, despite being evident, go unregistered from moment to moment—even features as obvious as the person you're talking to.

The phenomenon of inattentional blindness demonstrates one potential cause of this lack of sensitivity to features of the visual scene: We fail to register features of our experience that we're not paying attention to.

- If you think about this even briefly, it should be no surprise that this is the case. If you reject Cartesian mind-body dualism, then presumably you think that the existence of mind has an evolutionary explanation: Minds have survived as an evolved trait in humans either because having minds confers an evolutionary advantage or because having minds is a by-product of some other trait that confers an evolutionary advantage.

- In either case, it seems doubtful that being highly attuned to one's internal, subjective phenomenology would confer an evolutionary advantage. Imagine the caveman so captivated by the experience of color that he gets from gazing at a tropical flower that he is promptly eaten by the tiger he's failed to observe.

- Rather, what would confer an evolutionary advantage would be the ability to keep the world in mind: to pay attention not to your subjective experience, but to significant features of your environment. And though these phenomena from cognitive psychology call into question how well we pay attention to the qualitative features of our subjective consciousness, they don't call into question our ability to track features of our environment that we find significant.

- Of course, this is still bad news for Descartes's internalism. Descartes wanted to establish the entire edifice of knowledge on what he took to be our certain awareness of our own internal mental states. And what we've discovered so far suggests that our knowledge of those internal mental states is far from certain.

The failure of Descartes's internalist project needn't force us to abandon the project of putting knowledge on a secure footing. Rather, it just suggests that we need to find a different strategy. This is the project that we'll pursue in subsequent lectures.

READINGS

Chisholm, *Theory of Knowledge*.

Conee and Feldman, *Evidentialism*.

QUIZ

1 TRUE OR FALSE
According to Descartes, if you're being deceived by the evil demon, the demon could even make you believe that 2 + 2 = 5.

2 TRUE OR FALSE
According to Descartes, all of the evidence for knowledge must be internal to the mind.

3 TRUE OR FALSE
The door study, reported in 1998 by psychologists Daniel Simons and Daniel Levin, is an example of saccading.

4 Which of the following provides counterexamples to Descartes's claim that we have certain knowledge of our own occurrent visual experiences?

a Foveal vision
b Saccades
c Change blindness
d All of the above
e None of the above

Answer key can be found on page 207.

LECTURE 4

The Coherence Theory of Knowledge

Where Descartes suggested that all knowledge should be built up on secure foundations, his 19th-century opponents, along with their 20th-century descendants, rejected the idea that such foundations were possible—or necessary. In the place of such foundations, these opponents of Descartes's foundationalism suggested that the route to knowledge involves the way beliefs come together to form a system in which reasons are interconnected and mutually supporting. The idea was that knowledge has to do with the way beliefs cohere with each other. For this reason, this theory came to be known as coherentism.

INFALLIBILISM VERSUS FALLIBILISM

- The first problem that Descartes's successors raised for his theory of knowledge was his view that all knowledge must be certain knowledge. Descartes thought that this was true for two reasons.
  - He thought that the foundations of knowledge were facts that we could know with certainty. The paradigm example of this is the fact that you are now thinking. There is no way you can be mistaken about that fact at the very moment you are thinking it.
  - He thought that we build up our knowledge on our foundational beliefs by means of deductive reasoning.

- Deductive reasoning, however, is certain: If the premises of a deductive argument are true, then the conclusions arrived at on the basis of those premises must be true as well. This means that if we have certain foundations and further beliefs built up on those certain foundations by means of deductive reasoning, then we have certainty all up and down the edifice of knowledge! Because of this, Descartes's view is considered to be a version of infallibilism, the view that knowledge must be known with certainty.

- The problem with this is that it provides us with a very limited picture of the extent of knowledge. If knowledge requires certainty, then almost none of the facts that we take ourselves to know are facts that we actually know.

- For most philosophers, this seems like too drastic a step to take. This is why most philosophers—including almost all contemporary philosophers who study knowledge—are fallibilists. This means that they think that your knowing a fact is compatible with its being possible that you are wrong about that fact.

- Of course, because knowing something requires that what you know is true, knowing a fact isn't compatible with the fact's actually being wrong. It's just compatible with the possibility that the way you know the fact doesn't guarantee the truth of that fact with certainty.

- Many philosophers were convinced that Descartes's commitment to infallibilism was wrong. A number of them thought the answer was simply to abandon infallibilism but still maintain a commitment to Descartes's foundationalism. On this view, our knowledge requires foundations, but the way we build up our knowledge on those foundations doesn't guarantee certain truth.

Many philosophers, however, were distrustful of the idea that we could make sense of building up all knowledge on the basis of foundational beliefs. Because of this distrust, these philosophers rejected that the structure of knowledge ought to be thought of as a building, with the higher-level beliefs supported by the beliefs on the levels below. Instead, they suggested that we think of the structure of knowledge like a web—where each node of a web is mutually supporting and any tear in the web, at any point, can cause the whole structure to fail.

Philosophers who were distrustful of the idea that we could make sense of building up all knowledge on the basis of foundational beliefs noted, for example, that much of our perceptual knowledge is dependent on the other beliefs we have.

Without a knowledge of tools and their functions, for example, you wouldn't know that a piece of wood connected in a T shape to a piece of metal is a hammer. With that knowledge of tools and their functions, on the other hand, you immediately see a hammer and not a mere piece of metal and wood.

These philosophers—such as Willard Van Orman Quine, Donald Davidson, and Keith Lehrer—suggested that this weblike, mutually supporting structure of knowledge could be understood in terms of the notion of coherence. The view that the structure of knowledge should be understood in this way is known as the coherence theory, or coherentism.

COHERENTISM

- In one sense, the intuition behind coherentism is unquestionable. If one's beliefs are going to count as knowledge, it would seem that they must, at the very least, be consistent. And presumably coherence must, at the very least, include consistency.

- Think about the ways in which we provide positive arguments to others in favor of our theories: We do so on the basis of other characteristics of those theories that philosophers have counted as part of coherence. These include properties of theories such as predictive and explanatory power and simplicity.

- This gives us, then, at least four properties that belong to coherence: consistency, predictive power, explanatory power, and simplicity.
  - At a minimum, you would assume that a defensible belief would have to be consistent with the other beliefs that you have. You can't believe that it's snowing, for example, unless you also believe that it's cold enough to snow.
  - Of course, just having a consistent set of beliefs is a fairly weak standard to satisfy. This is where the other conditions come in. If you have a consistent set of beliefs that also does a good job of explaining the world, then that's even better. This is the point of including explanatory power as a requirement for coherence.
  - Even consistency and explanatory power, however, are too weak. Consider the example of ancient Greek mythology. The myths that the ancient Greeks told were consistent. Furthermore, they provided an explanation for the phenomena that they observed around them. If it thundered, it was because a particular god was angry. The problem, however, is that ancient Greek mythology is not very useful in making predictions.
  - Although explanations from Greek mythology might help explain why a certain event occurred, it wouldn't be very good in foretelling what events have yet to occur. Although Zeus's anger might be a good after-the-fact explanation for a thunderstorm, it is more useful to measure current temperature, wind direction, barometric pressure, and humidity if you're interested in forecasting what the weather will be like in the next few hours. In other words, a theory of weather that takes such facts into account will have greater predictive power than Greek mythology.
  - You might sometimes have two theories that are on equal footing with respect to their consistency, explanatory power, and predictive power. When that occurs, you

should choose the simpler of the two theories. You may have heard an expression that characterizes the intuition behind including simplicity as a criterion: "When hearing hoofbeats, think horses, not zebras."

- There are two small issues plaguing coherence theory. Neither of these issues by itself is devastating for coherence theory. They show that the theory needs work, but they don't show that it's unworkable.
  - The first problem for coherence theory is that the theory requires that you have consistent beliefs in order to have knowledge. The problem here is that nobody actually has completely consistent beliefs. In fact, it's common for people to hold inconsistent sets of beliefs, and if we require consistency for coherence—and coherence for knowledge—then we'd have the undesired result that nobody knows anything. Presumably, however, there may be some way to weaken the consistency requirement. What the coherence theorist would have to do is find some way to require a certain level of consistency but not perfect consistency. The first problem, then, says that the requirements of coherence theory are too strong.

Perhaps you believe that a politician that you admire is a good person, that nobody who cheats on his or her spouse is a good person, and that the politician you admire cheated on his or her spouse. That's an inconsistent set of beliefs.

  - The second problem says that they're unclear. Due to the notion that coherence involves the four criteria of consistency, explanatoriness, predictive power, and simplicity, there will be different ways to weigh the four criteria that will result in different theories. But in the case of knowledge, our theory should give us the one correct theory—and it's hard to see, given the fact that we can weigh the components of coherence differently, how we can do that with coherence theory.

- There are three arguments against coherence theory that raise serious difficulties for it: the problem of surprising evidence, the problem of experience, and the problem of truth.
  - For the coherentist, what makes one of your beliefs well supported enough to be a candidate for knowledge is for that belief to cohere with the other beliefs that you already have. One challenge for the view, then, is what to do about surprising beliefs. What it means to be a surprising belief is precisely that it doesn't fit well with the other beliefs that you already have.
  - The way to make sense of a surprising belief in a way that fits with coherence theory would be to explain away the belief so that it isn't surprising. The problem for the coherence theory, then, is that it suggests we should dismiss or explain away surprising experiences or evidence. As many cases from the history of science and technology demonstrate, this is often precisely the wrong attitude to take. We should embrace the surprising evidence. New discoveries and the development of new theories depend on our working on coming up with new explanations to make sense of the evidence that surprised us.
  - So, the problem of surprising evidence for coherence theory is that there's no room for unexpected evidence in a theory that makes coherence with existing beliefs the standard for knowledge. The problem, then, focuses on the coherence criterion for coherence theory.
  - The problem of experience targets a different aspect of the coherence theory. This problem has to do with the fact that the only elements that coherence theory evaluates are beliefs. According to the problem of experience, the exclusive focus on beliefs as the only source of evidence forces the coherence theorist to misunderstand the many cases in which it is experiences, rather than beliefs, that are the main source of evidence for knowledge.
  - Even more serious is the problem of truth. According to this problem, even the most coherent system of beliefs can still fail to achieve truth. If this is correct, however, then coherence cannot be the property of beliefs that, in optimal conditions, results in knowledge. Coherence alone—even a very high degree of coherence—is not sufficient to ensure truth. Thus, the theory is unworkable: One can have a highly coherent theory and have beliefs that cohere extremely well with that theory, but those beliefs can still be false.

Surprisingly, one of the clearest statements of the problem of truth was offered by a defender of coherence theory, the early 20th-century philosopher Brand Blanshard, who noted that the notion of coherence theory as a theory of justification stands in tension with a notion of objective truth.

> *Someone might hold that coherence with a set of beliefs is the test of truth but that truth consists in correspondence to objective facts. If, however, truth consists in correspondence to objective facts, coherence with a set of beliefs will not be a test of truth. This is the case since there is no guarantee that a perfectly coherent set of beliefs matches objective reality.*

Blanshard thought that, if faced with the choice between coherence theory and an objective notion of truth, we would choose coherence theory. However, abandoning an objective notion of truth is not a good strategy to pursue. Instead, we should reject coherence theory.

READINGS

BonJour, *The Structure of Empirical Knowledge*.

Lehrer, *Knowledge*.

QUIZ

1 TRUE OR FALSE
Descartes thought that some of our knowledge could be fallible.

2 Which of the following is not a property that belongs to coherence?
- a Consistency
- b Infallibility
- c Predictive power
- d Explanatory power

3 TRUE OR FALSE
That there is no single best way to balance the conflicting ideals of consistency, predictive power, explanatory power, and simplicity is a problem for coherentism.

4 Which of the following is a serious problem for coherence theory?
- a The problem of surprising evidence
- b The problem of experience
- c The problem of truth
- d All of the above
- e None of the above

Answer key can be found on page 207.

LECTURE 5

Externalist Theories of Knowledge

One of the ways in which philosophers attempted to deal with the difficulties faced by Descartes was to reject his foundationalism. The result was the theory of coherentism. However, coherentism faces its own challenges. There are two further challenges faced by both traditional foundationalist and coherentist theories that we have not yet considered. Thinking about these will help us formulate a new type of theory, one that wasn't anticipated until the beginning of the 20th century.

CHALLENGES FOR FOUNDATIONALISM AND COHERENTISM

- One of the further challenges for both traditional theories is that they ignore how a person's beliefs are actually formed and sustained. In philosophical terms, both are theories of propositional justification. This means that they only look at purely logical relations between a person's evidence and beliefs.

- Modern philosophers generally accept that, in order for you to be justified in your belief, it's not enough for you to have good evidence. Rather, your belief must be based on the good evidence that you have. The fact that you have that good evidence should explain why you hold the belief.

- Many philosophers characterize this idea by saying that theories of knowledge must include a basing relation; that is, adequate theories must include an account of what it is for beliefs to be properly based on the evidence so as to count as knowledge.

- Now consider the second problem facing both traditional foundationalism and coherentism. Both traditional foundationalism and coherentism are usually considered fallibilist theories in modern times. This means that they hold that you can know some fact even without certain evidence; that is, they say that you can know that fact even if your evidence doesn't give you an absolutely certain guarantee that it's true. This leaves these theories open to the challenge of so-called Gettier cases, named for philosopher Edmund Gettier.

- There are versions of both coherentism and traditional foundationalism that attempt to give an account of the basing relation and solve the Gettier problem. However, a new class of theories of knowledge emerged in the 20th century that offer an elegant way of achieving both of these goals at the same time. These are known as externalist theories.

PROCESS RELIABILISM

- There are now many versions of externalist theories. One of the simplest—known as process reliabilism—is both representative of the main features of externalist views and is capable of explaining many of the features of knowledge that we seek to explain.

A simple version of a Gettier case comes from Bertrand Russell:

- Suppose you're rushing to a meeting at noon in an unfamiliar building. You check the time on the clock in the hallway and it reads 11:55, so you believe that you have five minutes to get to the meeting location.
- Unbeknownst to you, however, the clock has stopped working hours before. Famously, though, even a stopped clock is right twice a day, and it just so happens that this clock stopped working exactly 12 hours ago, at 11:55 pm.

Here's the problem: You believe it's 11:55. It's true that it's 11:55. Your evidence for your belief is that the clock says it's 11:55, and for all you know, you have no reason to suspect that the clock isn't working. You're even basing your belief appropriately on your evidence.

Most philosophers, however, agree that you don't have knowledge in this case, because it's purely an accident that your evidence points you toward the truth. Had you looked at the clock a few minutes earlier or later, it would still have read 11:55 and your belief would be mistaken.

- One of the earliest statements of process reliabilism can be found in the work of mathematician F. P. Ramsey, who wrote in a 1929 paper,

> *I have always said that a belief was knowledge if it was (i) true, (ii) certain, and (iii) obtained by a reliable process.*

- With the exception of the inclusion of what Ramsey terms "certainty," this account of knowledge is very similar to ones that we would encounter today.

- If we reformulate the account to make it compatible with fallibilism, we're left with this: Knowledge is true belief on the basis of a reliable process. Here, the term "reliable" roughly means "consistent." And a "reliable process" is a process that reliably results in true beliefs.

Psychologists use the term "reliability" roughly to mean "consistency." In psychology, a measure is reliable if, given the same inputs, it'll yield the same outputs.

For psychologists, then, something could be reliable but wrong—as long as it's wrong consistently. For example, if your watch is running well but is set 23 minutes fast, then your watch is reliable, despite not being accurate.

When psychologists talk about the accuracy of a measure, they use the term "validity." So, for a psychologist, the best measures would be those that are both reliable and valid.

It's this that philosophers are aiming at with the notion of "reliable process."

- It's obvious how the process reliabilist theory can answer the first challenge—the challenge of explaining the basing relation. According to the process reliabilist, knowledge is true belief formed on the basis of a reliable process, a process that reliably yields true beliefs. And the answer to the question of how to determine the basing relation is baked right into the reliabilist theory: A belief that counts as knowledge will be appropriately based on the grounds that justify it because what it is for a belief to be justified is precisely to be caused by a process that reliably yields true beliefs.

- With regard to the second challenge—that of the Gettier cases—we need to consider the notion of "reliable process." Note that in order to think about reliability accurately, we'll have to suppose that we're considering a process to be reliable with respect to a certain environment. This is important, because processes that are reliable with respect to one environment may not be reliable with respect to a different environment.

- Consider a car whose engine is designed to function properly on generally well-paved roads in an industrialized nation. In a different environment—for example, in a desert with few roads—the engine might clog and become useless. The fact that the car's engine doesn't function reliably in desert environments does not mean that it isn't reliable, but rather that it's only reliable in the environments for which it was designed to function.

- So, our account of knowledge in terms of reliable processes would have to include a reference to the environment as well. We could, more properly, say something like this: Knowledge is true belief based on a process that is reliable (in the sort of environment in which that belief was produced).

- The challenge is to come up with an account of justified true belief that explains why Gettier cases—such as the stopped clock—are not cases of knowledge. And the process reliabilist theory allows us to do just this.

- Even though we have a true belief in the clock case, we don't have knowledge. The reason for this is that the process we've used—looking at a clock in the vicinity and relying on the information it conveys—is unreliable in an environment where the clock in the vicinity has stopped.

- For the process reliabilist, someone who believes on the basis of an unreliable process is not justified. We have a true belief, but no justification. And without justification, there is no knowledge. So, the process reliabilist can explain why Gettier cases aren't cases of knowledge.

Externalism, represented by process reliabilism, seems to enjoy two advantages over traditional foundationalism and coherentism. Unlike those theories, the process reliabilist has a simple answer to the challenges posed both by the basing relation and by Gettier cases.

EXTERNALIST VERSUS INTERNALIST THEORIES

- Despite the fact that they are rival views, foundationalism and coherentism have two key features in common.
  - The first feature is that what counts as support for beliefs are states that are internal to the believer. In the case of coherentism, the supports for your beliefs are just other beliefs and the way that the totality of your beliefs fit together to form a coherent whole. In the case of foundationalism, the support for your beliefs would include beliefs, but also your experiences, as well as the logic-based and content-based relations that connect them. The relations of logic would characterize the ways in which beliefs might be logically or inductively supported by other beliefs. The relations of content would characterize the ways in which your experiences support your beliefs.
  - Both theories can only conceive of the sorts of relations that provide support for beliefs in terms of logical, conceptual, or explanatory relations. This, then, is the second feature shared by both views. The ways that beliefs are supported are the sorts of ways that make sense upon reflection: deductive and inductive logic, conceptual connection, or coherence.
- Think of the justification that leads to knowledge as a process governed by two types of rules. The first type of rule—let's call them input rules—tells you what kinds of inputs you can introduce into the process. The second type of rule—let's call them transition rules—tells you what you can do to the inputs or to other intermediate results within the process. The output of the process will be justified beliefs.

> In the case of tax software, there are two kinds of inputs: one for the type of entry and the other for some number amount. The transition rules are the rules governing what effect those types of entries and number amounts have on the tax you must pay, and the output of the process is your tax owed or tax refund.

- With this way of thinking, we can characterize traditional coherentism as the view of justification that involves one type of input—beliefs—and one type of transition rule—coherence formation, however coherence theorists are able to make sense of that notion. Traditional foundationalism will have more inputs—beliefs, but also experiences of various types—and more transition rules—deductive inference, inference based on probability, and perhaps various other types of inference.

- The inputs may include not just the variety of inputs that the traditional foundationalist allows, but also nonmental states. The transition rules can also be broader than either coherentism or traditional foundationalism includes. They can encompass the sorts of transitions that foundationalists and coherentists describe—just as long as those transitions are in fact reliable processes in the environments where they're deployed—but they can also include additional reliable processes not countenanced by the coherentist or foundationalist.

- The input rules involved in both coherentism and foundationalism limit the inputs to justification to states that are internal to your mind—either beliefs alone or beliefs and other types of mental states, such as experiences. And the transition rules that coherentism and foundationalism allow are supposed to be the sorts of rules that you can recognize, simply through reflection, to be the sorts of rules that promote the discovery of truth.

- In contrast, the input rules in process reliabilism allow for states of your brain and nervous system, other states of your body, or even states of the world outside of your body potentially to serve as inputs to the processes that result in justified beliefs.

- And as for the transition rules, process reliabilism allows for the possibility that there might be reliable processes leading to the formation of justified beliefs that we cannot know in advance, through reflection alone; that is, for the process reliabilist, empirical research into the processes that humans (or nonhuman animals) use to discover true information about the world will be relevant in determining whether a certain way of forming beliefs should count as justification or not.

- This means that traditional foundationalism and coherentism are internalist because they stipulate that the basis for your beliefs must be internal to your mind and that the way your beliefs are justified must be something you can recognize, purely on the basis of reflection, as being a legitimate source of justification.

- Process reliabilism is externalist both because it allows that at least part of the basis of your beliefs may include states outside of your mind and because it allows for the possibility that you might not be able to recognize, on the basis of pure reflection alone, why the actual basis for your beliefs is in fact a legitimate source of justification.

Consider again that you see the clock read 11:55 and form the belief that you have five minutes to get to your meeting. Remember, the clock has stopped, though you don't know that it has stopped.

According to both the foundationalist and the coherentist, you are justified in this case. According to the process reliabilist, you are not.

- The advantage for process reliabilists is that they have a very simple explanation for why the Gettier case is not a case of knowledge: Although you, looking at the stopped clock, have a true belief in the case where the clock happens to be right, you don't know because you are not justified.

- The foundationalist and coherentist must give a more complicated explanation for why Gettier cases don't amount to knowledge. They've offered various explanations, all of which involve adding conditions to the analysis of knowledge. Rather than saying that knowledge is justified true belief and then explaining that justification is understood in terms of either foundationalism or coherentism, they must add an additional condition that explicitly rules out Gettier cases.

READINGS

BonJour and Sosa, *Epistemic Justification*.

Goldman, *Epistemology and Cognition*.

QUIZ

1 TRUE OR FALSE
The example of Jeeves the butler provides support for the necessity of a basing relation for knowledge.

2 TRUE OR FALSE
The stopped-clock example from Bertrand Russell is an example of a Gettier case that predates Edmund Gettier's famous article by almost 40 years.

3 TRUE OR FALSE
Reliability in psychology means exactly the same thing as it does for philosophers.

4 Which of the following is an externalist theory?

- **a** Descartes's foundationalism
- **b** The coherence theory
- **c** Process reliability
- **d** All of the above
- **e** None of the above

Answer key can be found on page 207.

LECTURE 6

Problems with Self-Knowledge

In traditional epistemology, self-knowledge is the most important form of knowledge. The "self" in self-knowledge refers to your mind—the states of your mind (the thoughts, feelings, emotions, and experiences) that Descartes wrote about as being the content of your mind. When they talk about self-knowledge, the epistemologists who think we have special knowledge of our own minds might be referring to transparency, which is the idea that your mind is transparent to you—that whatever is part of the content of your mind, you can know that it's part of the content of your mind. Or they might be referring to the infallibility thesis, which means that you're never wrong about the mental states you have.

TRANSPARENCY VERSUS INFALLIBILITY

Because, according to coherentism, the only way for a belief to be justified is to be part of a set of coherent beliefs, the coherentist will certainly not endorse infallibility about the beliefs we form about our mental states. However, at least for internalist coherentism, the transparency thesis is very important.

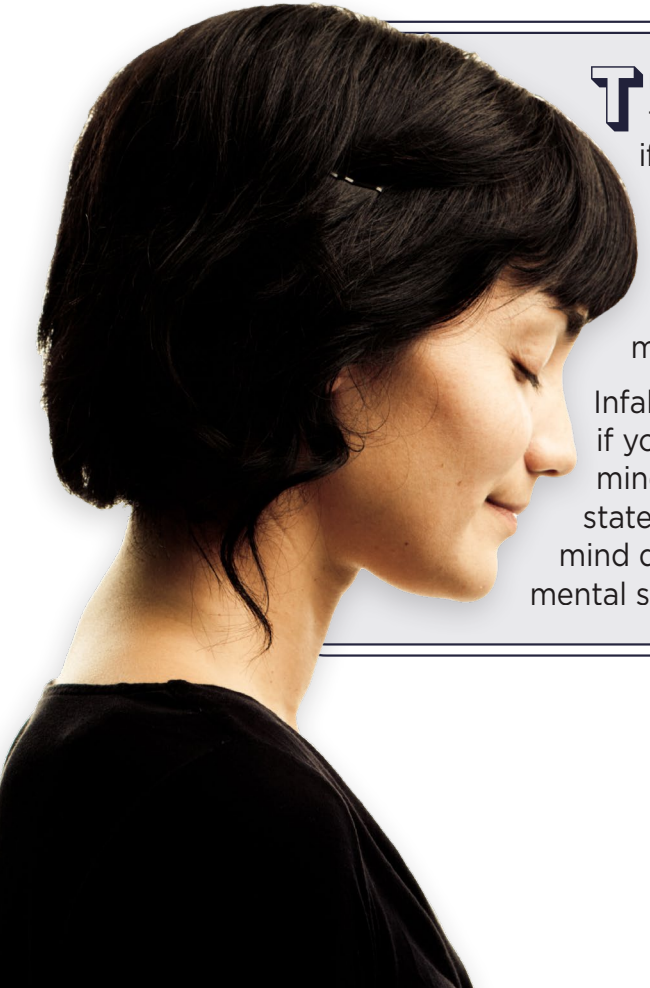

Transparency is the claim that if your mind contains mental state M, then you know that your mind contains mental state M.

Infallibility says that if you believe your mind contains mental state M, then your mind does contain mental state M.

When we talk about internalist coherentism, the "coherentism" part just refers to the fact that coherentists stress the importance of the coherence of your beliefs for your justification. Beliefs aren't justified individually but are justified by being members of sets of beliefs that cohere together, are useful in providing explanations and predictions, and are the simplest sets of beliefs that provide useful explanations and predictions.

The "internalist" part means two things.

- The bases for your justification are internal to your mind: They're mental states of yours. In the case of coherentism, they're beliefs.
- The ways you build up your justification are ways you can appreciate as contributing to your justification. In the case of coherentism, you can come to appreciate that you are justified, because you can, if you check, establish to your own satisfaction that your beliefs are in fact members of a coherent set of beliefs that is comparatively simple, explanatory, and useful for predictions.

If the coherentist is an internalist, then he or she thinks that to be justified you need to be able to check to see that your beliefs are part of a coherent set that is comparatively simple, explanatory, and useful for predictions. But that means that you need to have knowledge of what your beliefs are—you have to know what you believe.

- For any belief, if you believe it, then you know that you believe it—and this is just the transparency thesis applied to the case of beliefs. So, in order for coherence theory to work, the transparency thesis has to be true, at least in the case of beliefs.

- The coherentist wants us to survey our beliefs and see how they hang together. The better they fit, the more justified they are. In some ways, then, the coherentist thinks that searching for justification is like trying to tell a good story. You're looking for connections, trying to provide a picture that makes sense of all the component parts.

- Storytelling is an activity. It requires work on the part of the storyteller, looking for connections and drawing parallels. In contrast, the foundationalist doesn't think that our foundational justification requires much work on our part. In fact, the traditional foundationalist has a name for this foundational justification that makes the believer seem like just a passive recipient of his or her justification: the given, which refers to the sensory experiences that are just given to us when we engage with the world through our senses.

- Because of this, the foundationalist doesn't think that you need to survey your beliefs or other mental states in order to be justified in having those beliefs. When you're justified on the basis of experience, that's the most obvious thing in the world! Your sense experience, in such cases, just hands you your justification.

- For an experience to be foundational, it must provide you with very strong evidence for some belief. The strongest evidence it could provide, the traditional foundationalist suggests, is for a belief that you're having that sort of experience. And it's that belief about your experience that justifies your belief.

- But how much justification does your experience provide you for the belief that you're having that experience? When you believe that you have a headache, for example, how justified are you?

- If the infallibility claim is true, then if you believe you have a certain experience, then you really do have that experience. Your justification, in other words, is infallible. This is the sort of certain justification that Descartes was after.

- Of course, the certainty stops for you at your beliefs about your experiences. Those are certain, but what isn't certain is whatever you might infer from those beliefs.

- Regardless, however, what is certain is that the infallibility claim, at least with regard to our own phenomenal experiences, is of great significance for the traditional foundationalist. It is our supposed infallibility about our own experiences that provides

the secure foundation on which the traditional foundationalist thinks we must build up the entire structure of our justification and knowledge.

- An additional mental phenomenon to which the foundationalist might say we must have infallible access is the relations between our mental states and the more foundational states that support them. If you're going to be an internalist foundationalist and suggest that all your knowledge is based on a foundation of experiences—and beliefs about those experiences—to which you have privileged access, then presumably you'll also think that you have privileged access to the fact that your experiences support the beliefs about them.

THE TRANSPARENCY CLAIM

- When thinking about the transparency claim, the first thing that we need to do is try to consider only those mental states for which the claim is even remotely plausible. It should be pretty obvious that there are many aspects of our mind that aren't transparent to us at any given time. For example, your mind plausibly contains a whole host of memories that you're not even considering—and if you're not considering those memories, then you're also not currently believing that you have those memories.

> Philosophers call states of your mind that are occurring to you right now occurrent mental states. States that aren't occurrent are called dispositional states.

- So, the transparency claim only applies to states of your mind that are occurring to you right now. For example, if you now have a headache, then that is a state that is occurring to you right now.

- Let's assume that we're limiting the transparency claim to occurrent mental states, the states that are occurring to you right now. There is one further way that we can limit the transparency claim. Because the claim is important for coherentism but not for foundationalism, and because coherentists only care about the transparency claim with respect to beliefs, we can limit our discussion to the question of whether it is plausible that our beliefs are transparent to us.

- In other words, is it plausible to suppose that whenever you are currently believing something you are aware that you are believing it? That is what the limited transparency claim amounts to.

- Even if we limit the transparency claim in this way, it's still not likely to be plausible. The following are two reasons for this:
  - It's common for people to believe the opposite of what they know to be true. The starkest example of this is the phenomenon of denial—for example, when a cancer patient knows that he or she is facing death but refuses explicitly to admit it to him- or herself.
  - There are plenty of beliefs that you're unaware of at the moment at which you're having them. An instant's consideration would demonstrate that this is obvious. According to the transparency claim applied to occurrent beliefs, if you currently believe something, then you know that you believe it. Knowledge, however, is a form of belief. So, we would have to apply the transparency claim to that belief. So, if you currently know that you believe something, then you know you believe that you believe it—and so on. In other words, the transparency claim for beliefs means that any instant, you have an infinite number of occurrent beliefs. That's not a very palatable result!

THE INFALLIBILITY CLAIM

- There is good reason to doubt that your mental states—even the beliefs you're having right now—are all transparent to you. But suppose you're attending to an experience of yours and form a belief about it. Isn't it plausible to suppose that any belief about an experience to which you're now attending would have to be true? How could you be wrong in a belief about your own current experiences?

- This is the intuition behind the internalist foundationalist endorsement of the infallibility claim, at least with respect to occurrent experiences. And the internalist foundationalist needs the infallibility claim to be true in order to argue that our current experiences form an infallible foundation on which to build up the rest of our knowledge.

- Unfortunately for the internalist foundationalist, however, there are good reasons to question the infallibility claim, even when applied to our current experiences. Two such reasons come from philosophical considerations and psychological studies.

The Philosophical Argument

- Although the spectrum of color, from red to violet, is effectively continuous, humans' ability to distinguish colors from one another is not. This means that if the range of colors is depicted as a gradient, observers will perceive two different colors that are sufficiently close to one another on the gradient as being the same color.

- Consider three colors: The first color and the third color are as close as possible on the gradient while still being distinguishable, and the second color is halfway between the first color and the third color. Suppose that because the second color is so close to the first color, it's indistinguishable from the first color. Similarly, because the second color is so close to the third color, it is also indistinguishable from the third color.

- It would seem that the following things are true of you if you're observing the three colors:
  - You have the experience that color 1 is the same as color 2.
  - You have the experience that color 2 is the same as color 3.
  - You have the experience that color 1 is different from color 3.

- So, experiencing two colors as being the same is not a transitive relation—one like identity in math, where if $a = b$ and $b = c$, then $a = c$. But the problem is that sameness of color is a transitive relation. If color *a* is the same as color *b* and color *b* is the same as color *c*, then color *a* must be the same as color *c*.

- One way to deal with this problem is to suggest that the properties of experiences are somehow indeterminate, so that even though transitivity applies to sameness of color, it doesn't apply to experiences of sameness of color. But let's focus on the other way to deal with the problem, which is to deny that beliefs about experiences are infallible. In other words, when you believe you're experiencing two colors as the same, there's no guarantee that this really is the content of your experience.

The Psychological Argument

- There is a flood of evidence from cognitive psychology and brain science that substantiates the claim that we are often mistaken about the character of our experiences and the causes of those experiences and that we are unaware of the prevalence of those mistaken beliefs. The widely studied placebo effect is one example.
- Despite the wealth of such studies, our own subjective experience is that we are good judges of the subjective qualities of our experiences—and of the ways in which we base our judgments on those experiences. Yet if the arguments described here are correct, we have good reason to question that subjective experience about our own infallibility.

READINGS

Gertler, *Self-Knowledge*.

Hurlburt and Schwitzgebel, *Describing Inner Experience?*

QUIZ

1 TRUE OR FALSE
Transparency is important for coherence theory but not for foundationalism.

2 TRUE OR FALSE
Infallibility is important for coherence theory but not for foundationalism.

3 Which of the following is called into question by the phenomenon of denial?

a Transparency
b Infallibility

4 Which of the following is called into question by the color gradient and the just-noticeable-difference threshold?

a Transparency
b Infallibility

5 Which of the following is called into question by psychological studies that provide evidence against our reliability in judging the subjective qualities of our experiences?

a Transparency
b Infallibility

Answer key can be found on page 207.

LECTURE 7

Does Sense Perception Support Knowledge?

One of the most significant sources of our knowledge—and particularly our knowledge of the world—is sense perception. One of the oldest debates in philosophy involves the implications of our reliance on perception for our contact. Do the senses serve as an obstacle, standing between us and the world? This way of thinking about the role of the senses is misleading. Worse, this way of thinking is ultimately irrelevant to the question of whether—and how—the beliefs that we form on the basis of perception are justified. However, this way of thinking about the senses is common, even among scientists who study sense perception.

In *Making Up the Mind: How the Brain Creates Our Mental World*, noted British psychologist and neuroimaging pioneer Chris Frith explains:

> *Even if all our senses are intact and our brain is functioning normally, we do not have direct access to the physical world. It may feel as if we have direct access, but this is an illusion created by our brain.*

INDIRECT PERCEPTUAL REALISM

- In philosophy, the view that you only ever perceive the world indirectly is called indirect perceptual realism. It's realism because the view doesn't deny that there really are objects in the world that we're perceiving; it just holds that whenever we perceive those objects in the world, it's through the intermediary step of perceiving our experiences.
- One further argument one might point to in order to support the idea that our senses create our mental world is this: Different animals have different perceptual abilities. For example, bats and dolphins use sonar to navigate and to locate and track prey. Given this, it's not implausible to suppose that different animals have experiences that feel very different.

Philosopher Thomas Nagel has a celebrated essay entitled "What Is It like to Be a Bat?" in which he argues that the radically different way in which the bat takes in information about the world through sonar means that we humans actually have no way of knowing what it's like to be a bat.

- It's easy, then, to appreciate the popularity of the view that the senses are an obstacle between us and the world—or, at the very least, to appreciate why you might see the senses as potentially unreliable informants. The problem with all of these arguments, however, is that they completely misrepresent what sense perception does.

- The arguments all share two features: They suggest that what sense perception really gives us is a feeling or experience with a particular phenomenal quality and that this feeling or experience is what serves as evidence for some further belief about the world. But these phenomenal qualities are in our minds. They're all feelings in the mind, even if they might seem to be located somewhere outside of our heads.

- Furthermore, the arguments also suggest that, because the feeling or experience is obviously distinct from the world, this introduces the possibility that the feeling or experience could actually be misleading us about the true nature of the world. In fact, part of the way that our senses mislead us is precisely by making us think that what is in our heads is actually in the world.

- The first of these features is more fundamental, then. It's because what sense perception immediately gives us is an experience with its own particular phenomenal quality—and because this experience is distinct from the world beyond it—that there is the possibility that the experience could be misleading us about the true nature of the world.

- So, the indirect perceptual realist picture of perception is one in which the primary function of the senses is to provide us with experiences with their own unique phenomenal qualities. Then, on the basis of these experiences, we infer the existence of objects in the world that correspond to those experiences or, at the very least, cause them.

- There are two main problems with this picture of perception—problems so central that they make it difficult to see how the indirect perceptual realist picture could be at all plausible.

- The first problem is that this picture of perception, according to which we infer the existence of objects in the world based on our sense experiences, doesn't fit with our understanding of how the brain actually processes perceptual information.

- There are many stages of perceptual information processing in the brain before that information reaches a stage at which it can plausibly be called an experience. By the time it reaches that stage, however, it is no longer the sort of purely phenomenal experience that the indirect

perceptual realist is talking about. Instead, the sort of experience that results from those many earlier stages of brain processing already contains a great many assumptions about external world objects.

It's hard to imagine a more phenomenal kind of experience than a pain experience. But even a pain experience cannot be understood as a pure phenomenon.

For one thing, you don't experience a pain as being located somewhere in your mind. Rather, you experience a pain as located somewhere in physical space—it's a headache or backache.

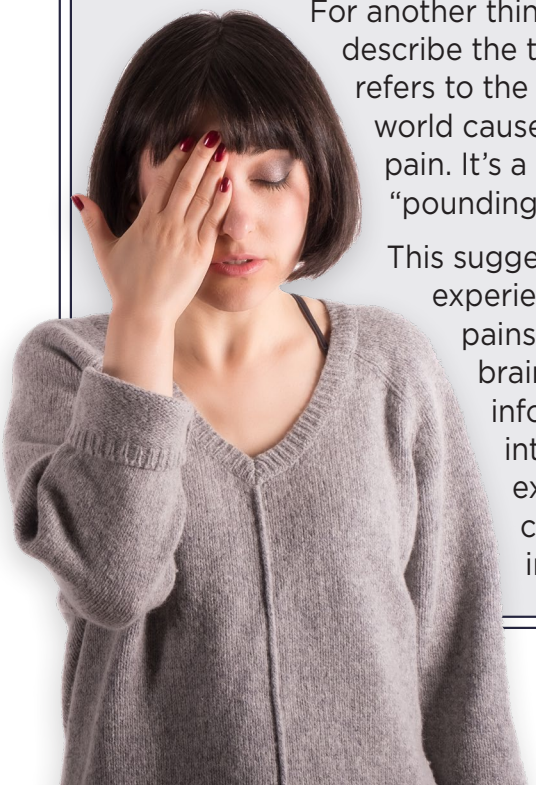

For another thing, even the way we describe the type of pain generally refers to the types of external world causes that lead to the pain. It's a "stabbing" pain or a "pounding" pain.

This suggests that even with experiences as raw as pains, by the time your brain has processed the information that goes into the experience, the experience itself already contains a great deal of information.

This is obvious in the case of sense experience. For healthy perceivers, our visual experience, for example, comes to us with a great deal of information already processed and ready for use, no inference on our part required.

Consider the case of perceivers who are not healthy. Certain brain injuries, resulting from tumors or trauma, can affect centers of the brain responsible for processing visual information and categorizing it according to the objects it contains. People with damage to these areas are said to suffer from visual agnosia.

Here's one of the scenarios that researchers who study visual agnosia have observed.

First, the researchers present the sufferers of agnosia with pictures of common objects, such as pens, and ask them to sketch what they see. They're able to do so with no difficulty. Then, researchers ask them to label the objects that they've drawn. The agnosics are unable to do this.

Furthermore, if you ask agnosics to define those objects—"What are pens for?"—they are able to do that.

If perception simply involved inference, then this condition shouldn't exist. Agnosics should easily be able to infer that the experiences they have correspond to these common, everyday objects and then to name those objects they see. Because they cannot do this, this suggests that the brain processes visual information prior to experience—so that healthy perceivers already enjoy experiences of the objects in their environment as experiences of those objects.

- The second problem with indirect perceptual realism is that it sees the role of the senses as being able to provide us with phenomenal sensory experiences. But this is simply mistaken. Rather, the role of the senses is to provide us with information about the world around us.

- These two problems may seem bad enough for indirect perceptual realism, but there is an even bigger problem for the view, at least from our perspective: According to indirect perceptual realism, all we immediately perceive is our own perceptual experiences, rather than the objects that those experiences are about. But that's a claim about the immediate *objects* of our experiences.

- The technical term for that subject matter is metaphysics—the study of the basic building blocks of reality and their relations to each other.

That's not what we're primarily concerned with. We're interested in *knowledge* and how to get it. So, we're interested in the building blocks of knowledge, such as beliefs, and what makes some beliefs count as knowledge.

- And the point now is that the beliefs we form on the basis of sense perception are not primarily beliefs about our experiences. They're beliefs about phenomena in the world that we care about: people, places, and things.

- But indirect perceptual realism is a theory about the objects of our perceptual experiences, not about the objects of the beliefs based on those experiences. So, even if indirect perceptual realism were true, that would have no direct bearing on the relation between sense perception and knowledge.

SENSE PERCEPTION AND KNOWLEDGE

- The biggest problem with indirect perceptual realism is that the theory is simply irrelevant to perceptual knowledge. So, the question that we should focus on isn't whether indirect perceptual realism is true but, rather, how sense perception supports knowledge. When we know something on the basis of what we perceive, how is our knowledge based on that perception?

- There are different ways to answer this question. One is to ask whether the relation between sense perception and knowledge supports a foundationalist or a coherentist theory. The other is to ask whether the relation between our senses and our knowledge supports internalism or externalism.

In the remainder of this lecture, we'll begin to tackle the first question. In particular, we'll consider whether there is a good reason to suppose that coherentism offers the best explanation for how sense perception supports knowledge.

In the next lecture, we'll first consider whether there is good reason for thinking that internalist foundationalism offers the best explanation for how sense perception supports knowledge or whether we ought to endorse some form of externalist explanation instead.

- There are two different types of distinctions that we can make when we talk about perceptual processes, and the relation between those two types of distinctions is liable to cause confusion.

- First, it's important to distinguish between computations that you perform and computations that your brain performs but that you have no access to. Philosophers sometimes call the first type of processes—the ones that you normally have access to and at least some control over—personal processes. And they call the second type of processes—the ones that are performed by organs or systems of organs in the body, without direct access or control by the person whose body it is—subpersonal ones.

- The second distinction is between what cognitive psychologists refer to as top-down and bottom-up processes.

- A bottom-up process is one that arrives at a result from more basic inputs, where each of those inputs is computed independently. A bottom-up theory of visual perception would suggest that information is transmitted in one direction from the retina to the visual cortex. On this type of theory, each stage in the mechanism for the processing of visual information involves increasingly complex analysis of the information transmitted from the immediately preceding stage.

- Common versions of bottom-up theories include the idea that the neurons involved in each stage of visual information processing are highly specialized to perform a very specific type of computation. For example, there are neurons that just perform detection of light versus dark, or of edges that are oriented in a certain direction, or of motion.

- In opposition to bottom-up views of perceptual information processing are top-down views. According to top-down theories, the brain is able to use higher-level information as feedback for more basic sensory processing mechanisms in the brain. Top-down theorists reject the idea that the processing of sensory information in the brain moves in only one direction, from the more basic to the more complex. Instead, they suggest that the more complex computational mechanisms in the brain can provide feedback for earlier mechanisms.

- A number of the pieces of evidence cited in favor of top-down theories have to do with the fact that your semantic knowledge—your knowledge of concepts and meaning—can affect what you perceive.

- There seems to be at least some evidence for top-down processing of sensory information in the brain. And top-down processing involves taking context into account before deciding how to interpret information. And this, in turn, should remind you of coherentism—about the way in which beliefs can be justified enough to count as knowledge.

- Remember, coherentism says that what makes an instance of true belief count as knowledge is that the belief fits together with the rest of your information. And the top-down theory says that your sensory processing system takes different pieces of information and interprets them based on how they best fit together given the context.

- Does that mean that the cognitive science of perception supports coherentism? That would be very strong support for the theory, but unfortunately for coherentism, top-down theories don't offer support for any theory of the structure of knowledge.

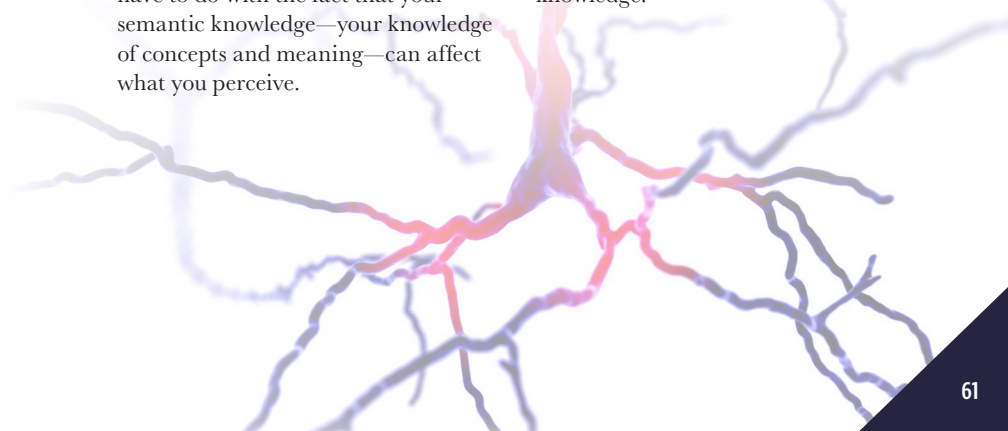

If all sense perception involved top-down processing, then we would expect newborns not to be able to perceive anything at all until they developed the concepts to help them make sense of what William James called the "blooming, buzzing confusion" of their immediate sensory stimuli.

This, however, is not the case. For example, newborn human infants seem to show a preference for face shapes over scrambled stimuli—as early as a few minutes after they're born.

- One problem is that top-down theories cannot tell the whole story of sense perception. There must be at least some room for bottom-up explanations as well.

- The second problem for coherentism is even more serious. In order for some empirical evidence to support one of the theories of internalist justification, such as coherentism, it would have to involve personal computational mechanisms. That's because coherentism is a theory about what goes on in the mind, rather than about the subpersonal mechanisms that are plugging away below the surface of your awareness or control.

- The debate between top-down and bottom-up explanations, however, is best understood as involving just those sorts of subpersonal mechanisms. That's because the debate is supposed to involve the sorts of computational processes to which you normally have no access and over which you normally have no control. So, even if the best explanation of sense perception is top-down—and there is reason to question whether that's even the case—it would still not support coherentism.

READINGS

Lehrer, "Coherence, Truth, and Undefeated Justification."

Russell, *The Problems of Philosophy*.

QUIZ

1 TRUE OR FALSE
The argument from illusion attempts to establish that what you immediately perceive in sense perception is your own subjective experiences.

2 TRUE OR FALSE
People with visual agnosia can accurately sketch the features of the objects that they see, but they cannot recognize those objects.

3 TRUE OR FALSE
The most plausible understanding of our sensory systems is that their role is to provide us with phenomenal sensory experiences that enrich our inner mental lives.

4 What does the discussion of the *Citizen Kane* example tell us?

a Indirect perceptual realism really is the only correct view of sense perception.

b Indirect perceptual realism only makes sense in the case of viewing movies, photographs, etc.

c Even if indirect perceptual realism is true as a theory in metaphysics, it is in fact irrelevant to discussions in epistemology.

5 David Hubel and Torsten Wiesel's research establishing the existence of specialized neurons for the detection of very specific features of visual stimuli would seem to be an example of which of the following?

a Subpersonal, top-down processing

b Subpersonal, bottom-up processing

c Personal, top-down processing

d Personal, bottom-up processing

Answer key can be found on page 207.

LECTURE 8

Perception: Foundationalism and Externalism

In the previous lecture, we ruled out the internalist theory of coherentism to provide the best explanation of the support for knowledge based on sense perception. This leaves us with one more type of internalist theory to consider: foundationalism. There are two types of foundationalism to evaluate: experience-based internalist foundationalism and naive foundationalism. Experience-based internalist foundationalism isn't very plausible; in fact, it suffers from many of the same problems that plagued indirect perceptual realism. And naive foundationalism doesn't provide us with an explanation of why sense perception provides evidence for knowledge; it doesn't tell us, in other words, why finding out that someone believes some event occurred because he or she saw it him- or herself makes us more confident that the person has knowledge that the event occurred. If the internalist theories

don't work, that leaves us with externalism, and the evidence from the cognitive psychology of perception supports the idea that an externalist theory of how perception contributes to knowledge is indeed the most plausible.

EXPERIENCE-BASED INTERNALIST FOUNDATIONALISM

- Experience-based internalist foundationalism is the view that all of our beliefs on the basis of sense perception are actually based on beliefs about the immediate qualities of our experiences—experiences of things like color patches and shapes—rather than experiences of the objects that produce the stimulations of our sense organs that result in perception.
- There are at least two big reasons why an experience-based internalist foundationalism won't work.

Both reasons why an experience-based internalist foundationalism won't work are related to points made in the previous lecture about what's wrong about indirect perceptual realism.

- It suffers from the same problem as internalist coherentism when it comes to the issue of where computation, or conscious inference, takes place in the process from perceptual stimuli in the sense organs to perception-based knowledge.
  - On the experience-based internalist foundationalist picture, our knowledge on the basis of perception is ultimately founded on our knowledge of more basic facts about our experience. So, every instance of perception actually involves an inference—an inference that, if we pay close enough attention to our own mental processes, we can be aware of. But this idea is simply absurd.

- As explained in the previous lecture, there is in fact a great deal of computation that occurs at the early stages of perception. And we can call these different types of computation "inferences" if we wish, but it is simply a mistake to think that we can access those "inferences," however closely we attend to our mental processes.
- Much of the computation—or inference, if you'd like—in perception happens prior to our awareness. Much of that computation is something your *brain* does, not something *you* do.

- It is also wrong about the content, or nature, of our experiences in perception.
  - According to the standard experience-based foundationalist picture, our experiences are the basis of our inferences about the objects in the world that cause those experiences, and those inferences then form our knowledge about the world. We've already seen reasons why this picture is inaccurate in the previous lecture, when we discussed the various problems faced by indirect perceptual realism.
  - Experience-based foundationalists like to act as if their position is more scientific and that it better accounts for the way experience involves interpretation on the part of the perceiver. In order for experience-based foundationalists to be truly scientific, however, they must be able to specify exactly what the basic building blocks of our experiences are—the colors, shapes, and so on that make up those foundational experiences on the basis of which we construct everything else.
  - The problem is that the experience-based foundationalist can't do this. There is the fact about what sort of information hits the retina and the fact about how that information is processed at various stages of the perceptual process, but those different types of information aren't the experiences that the foundationalist is talking about. This means that the experience-based foundationalist story is actually just that—a story, or fairy tale, with no basis in science. Science tells us that the information recorded by the retina, to take one example, is nothing like the sort of visual scene that even plausibly describes your experience. That visual scene is the result of a number of computational processes in the brain, none of which are accessible to awareness.

NAIVE FOUNDATIONALISM

- Why not just abandon the part of experience-based foundationalism that we find objectionable? That is the idea that the ultimate basis of our perceptual knowledge is our experiences, described as highly abstract sensory qualities like color, shape, and so on. If we abandon that objectionable part of the experience-based foundationalist view, what we're left with is naive foundationalism.

- Naive foundationalism rejects the idea that we must ground our perceptual knowledge in beliefs about the qualities of our experiences. It involves the claim that we can base our perceptual knowledge simply on the perception of the objects in the world that our knowledge concerns.

- According to the naive foundationalist, the basis of your perceptual knowledge that something is the case just *is* the perception that it's the case. The problem with this view is that sometimes it seems to you that you perceive something and it turns out that you're relying on a misleading perception. On the naive foundationalist view, there's nothing to explain the difference in the support you have when you misleadingly perceive something to be the case and when you correctly perceive something to be the case.

- That's why we can't simply accept naive foundationalism. Externalism will allow us to explain why certain perceptions support knowledge whereas other, misleading perceptions don't. The difference the externalist can point to—but the naive foundationalist can't—is the reliable accuracy of the perceptual systems that lead to knowledge.

- Whether a perceptual system is reliably accurate or not is a fact external to the mind of any particular believer. And the naive foundationalist can't appeal to reliable accuracy because, as an internalist, he or she can't appeal to any criteria that are outside of the mind of the believer.

If the information presented in the previous lecture and this lecture is correct, neither coherentism nor foundationalism does a good job of explaining how sense perception gives us evidence for knowledge.

EXTERNALISM

- Two aspects are essential to externalist theories.
  - They make reliable accuracy the central quality for determining whether a certain cognitive process is a source of knowledge. In the case of sense perception, this means that only those sensory processes that reliably give us accurate information about the world will count as potential sources of knowledge.
  - They don't require us to be aware of the operations of the processes that give us information. This means that externalist theories—in contrast with internalist theories—allow that it is possible that we can be unaware of the process we are employing to acquire information and of whether that process is in fact reliably accurate. In the case of sense perception, this means that we at least sometimes acquire information on the basis of sensory processes of which we're unaware or about whose accuracy we're unaware.
- There is a great deal of support from cognitive psychology and neurobiology for both of these components of externalism. One example is the existence of dual streams of cognition for the processing of visual information, and another involves what neurologists call brain bilateralization, the fact that the two halves of the brain are specialized for different tasks that they can perform independently.
  - One way that brain functions are specialized is between the dorsal and ventral streams. The visual cortex, the region of the brain that is responsible for receiving information from the eyes, is located at the back of the brain. The dorsal stream—or the "where" pathway—runs along the top of the cortex. It's the visual system that aids in determining the spatial location of the objects around us, preparing us for any actions that we might want to perform with those objects. The other pathway is called the ventral stream—or the "what" pathway—and it is located on the underside of the cortex. It's the network responsible for the visual identification of objects.

Evidence for the dual streams of visual information processing comes from studies conducted by neuroscientists David Milner and Melvyn Goodale.

- The brain has two halves, or hemispheres, which are joined by a very densely packed strand of nerve cells called the corpus callosum. There are some neurosurgical procedures, called split-brain surgeries, that require cutting the corpus callosum so that the two hemispheres of the brain can no longer communicate. For many years after the first split-brain surgeries were performed, neurosurgeons didn't think that cutting the corpus callosum had any effect on brain function. Certainly, split-brain patients didn't claim to notice any difference before or after the surgery. That was the thought until neuroscientist Michael Gazzaniga began studying split-brain patients and found that the two sides of the brain are unable to make sense of what the opposite side is aware of.

- The phenomena of dual-stream pathways for visual information processing and of split-brain cases seem to provide evidence for both of the aspects of sense perception that support externalism.
  - Both involve evidence that some cognitive systems are more reliably accurate in either gathering or evaluating information in perception.
  - Both support the idea that we are often unaware of the processes we actually employ in sense perception and that we are often equally unaware of whether or why the processes we employ are reliably accurate.

- This suggests that, at least in the case of sense perception, externalism provides the best account for explaining the structure of knowledge.

To study the inability of the two sides of the brain to make sense of what the opposite side is aware of, Michael Gazzaniga and his team of researchers flashed an image to each eye of a split-brain patient while keeping that image hidden from the other eye. At the same time, they gave each patient's hand an option to choose from one of four pictures that best matched the image. In other words, the left hand could choose from the pictures that matched the image that the left eye saw, and the right hand could choose from the pictures that matched the image that the right eye saw.

One very dramatic illustration of the results of Gazzaniga's research with split-brain patients is the case of a subject referred to as P.S.:

> *P.S. had seen a chicken claw with his left brain, and his right hand chose a picture of a chicken. His right brain had seen a picture of a snow scene, and his left hand picked up a picture of a shovel. When asked why he had done all of this, he said from his left hemisphere, 'The chicken claw goes with the chicken, and you need a shovel to clean out the chicken shed.'*

This suggested to Gazzaniga that there is a special mechanism in the brain—which he calls the interpreter—that is responsible for explaining our behavior to ourselves. It concocts a plausible story for us about what is going on in our heads when we pause to think about what we're perceiving or thinking. But importantly, what the interpreter concocts is just a story—one that's often not very accurate.

READINGS

Jackson, *Perception*.

Noe, *Action in Perception*.

QUIZ

1 What is the type of foundationalism that claims that the basis of sensory knowledge is our immediate, qualitative experiences of phenomena like color patches, shapes, or regions of light?

a Experience-based internalist foundationalism

b Naive internalist foundationalism

2 TRUE OR FALSE

Patrick Cavanagh's work on the physics of shadows provides evidence in favor of the claim that the inferences on which our sensory knowledge is based are the result of the sort of personal processes that the experience-based internalist foundationalist suggests.

3 TRUE OR FALSE

The problem for naive internalist foundationalism is that the naive foundationalist cannot appeal to the reliable accuracy of sense perception to explain why sense perception supports knowledge.

4 If you see an object in your vicinity and track its location relative to you, that's a result of what?

a Your brain's dorsal stream

b Your brain's ventral stream

5 According to Michael Gazzaniga, our conscious experiences are a result of which of the following?

a The intrinsic experiential qualities in the objects themselves

b A special module in the left hemisphere of the brain that makes up an explanation of the information processed in the brain

c Brain activities distributed across both brain hemispheres

d All of the above

e None of the above

Answer key can be found on page 207.

LECTURE 9

The Importance of Memory for Knowledge

Our faculty of memory is central to our knowledge, and philosophers have had a number of things to say about the topic. Some of the philosophical theories are specific to the topic of memory while others are more related to the kinds of questions we have considered across a number of lectures. This lecture will address questions raised by the topic of memory specifically, including questions about what the objects of memories are. Are they stored experiences in the mind, or are they past events themselves? The questions also include whether memory merely preserves beliefs and knowledge that we've already acquired or whether we can ever gain new knowledge on the basis of memory.

TYPES OF MEMORY

- There are three main types of memory: long-term, short-term, and working memory.
  - Short-term memory stores small pieces of information for no more than a few seconds or so.
  - Working memory, like short-term memory, also has to do with storing information for very short periods of time to aid with both reasoning and action.
  - The type of memory that philosophers are concerned primarily with is long-term memory.
- There are three different types of long-term memory, roughly based on the type of information stored: procedural, semantic or declarative, and episodic.
  - Procedural memory supports the ability to perform skilled actions. For example, if you know how to play a certain musical piece on the piano, that skill relies on procedural memory.
  - Semantic or declarative memory involves the memory of events in terms of their descriptions. It involves information stored with the aid of language.
  - Episodic memory involves the memory of experiences. Instead of being stored with the aid of language, episodic memory is what allows us to replay previous experiences at a later time.

Though psychologists nowadays generally distinguish the concept of working memory from that of short-term memory, those differences don't concern us here.

The summary of the reason for distinguishing between them is that short-term memory was initially thought to be a single, undifferentiated short-term information storage mechanism, whereas theories of working memory generally involve more specialized types of short-term storage for accomplishing different cognitive tasks.

Interestingly, the three types of long-term memory seem to be distinct, in that you can have one type without having either of the other two.

A phenomenon in cognitive psychology called aphantasia involves having little or no ability to form or manipulate mental images—no vivid imaginations or daydreams, for example. One of the phenomena often associated with aphantasia is severely deficient autobiographical memory, which is basically a lack of episodic memory capacity.

EPISODIC MEMORY

- Episodic memory—memory of previously experienced objects and events—forms the basis of central uses of memory, such as eyewitness accounts and autobiographical memory. For that reason, it has been the most-discussed type of memory among philosophers.

- What are the contents of our episodic memories? Are they the current memory experiences, or are they the past objects and events that we experienced? Do our memories themselves exist wholly in the present, or are they in part dependent on past objects or events?

- Bertrand Russell believed that he had an argument to demonstrate that our memories are themselves not dependent on the past but are simply stored experiences:

> *There is no logical impossibility in the hypothesis that the world sprang into being five minutes ago, exactly as it then was, with a population that "remembered" a wholly unreal past. There is no logically necessary connection between events at different times; therefore, nothing that is happening now or will happen in the future can disprove the hypothesis that the world began five minutes ago.*

Russell takes from this argument that because there is no logical impossibility in imagining a scenario in which all of the mental states that we take to be memories might actually refer to a supposed past that never took place, we should not define actual memories in terms of the past either. For Russell, the difference between our memory states and mere imagination lies solely in the fact that, in the case of memories, we *believe* them to be about the past.

Summing up, Russell suggests:

Memory demands (a) an image, (b) a belief in the past existence .... The believing is a specific feeling or sensation or complex of sensations, different from expectation or bare assent in a way that makes the belief refer to the past; the reference to the past lies in the belief-feeling, not in the content believed.

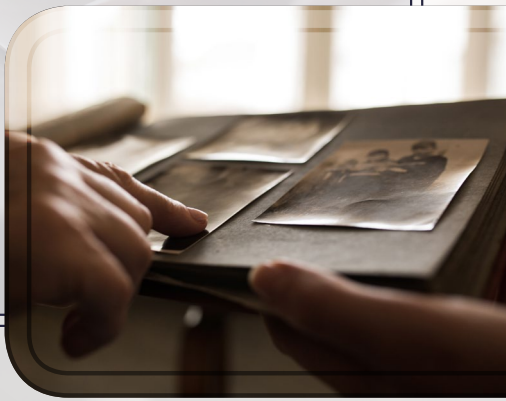

- There's a fatal problem with Russell's view, however. The problem is that Russell's view actually relies on the traditional view of memory that he claims to reject.

- Scottish philosopher Thomas Reid made this point very forcefully. He was arguing specifically against his contemporary David Hume's theory of memory, which was a forerunner of Russell's.

- Reid first notes that the commonsense view of memory is that it is an immediate knowledge of something that occurred in the past. And he notes that the view he's arguing against—the view held by both Russell and Hume—is that there isn't any such knowledge in the mind.

- That's because, according to both Russell and Hume, "memory is nothing but a present idea or impression." In contrast to Russell's view, for Thomas Reid the object of our memories are the past events themselves that we remember. In fact, as Reid has argued, this is precisely what distinguishes memory from perception or mere imagination.

- Because your current experiences don't belong to the past but rather are currently occurring, Reid argues in particular that the objects of your memories cannot be your current experiences. If Reid is right about this, then the Russell-Hume view of memory is wrong. Memories are not simply a type of current experience. Instead, they essentially involve a link to some previous object or event.

THE CAUSAL THEORY OF MEMORY

- Despite the fact that Thomas Reid offered a devastating criticism of Russell's view of memory almost 150 years before Russell proposed his theory, it wasn't until the middle of the 20th century that philosophers began to appreciate Reid's idea that memory requires a connection to a past object or event.

- Then, in a 1966 article entitled "Remembering," C. B. Martin and Max Deutscher provided an argument in favor of what is known as the causal theory of memory. According to that theory, memory involves a certain type of causal connection between your current cognitive state and the object or event that you remember. Since the article's publication, the causal theory has become widely accepted within the philosophical community.

In a 2017 article in the *Stanford Encyclopedia of Philosophy*, philosophers of memory Kourken Michaelian and John Sutton wrote,

> *The idea that remembering is characterized by an appropriate causal connection has … taken on the status of philosophical common sense.*

- Although the broad idea of some sort of connection between the past event and your current cognitive state is important for memory, recent discussions in philosophy have continued to pick away at specific components of the causal account.

- Consider a case in which it seems to you as if you remember a certain object or event. What makes that apparent memory a genuine instance of memory is if your seeming memory has a reliably stable, continuous brain-based connection to the past object or event that you seem to remember. This "reliably stable connection" view has two advantages that it inherits from the classic Martin and Deutscher account.

- First, it allows us to distinguish a number of cases of merely apparent memories from those of genuine memories.

- Suppose Harry visited Disneyland when he was four with his family. While there, he got separated from his family. Years pass and Harry forgets that experience. One day, Harry reads a book about a character who, at the age of four, visits Disneyland with his family and gets separated from them. The story makes an impression on Harry, so much so that years later, he comes to take the memory of the experience of the character in the book as if it's his own. When Harry later remembers the experience of getting lost at the age of four at Disneyland, it's because he read it in the book—not because it really happened to him. In a case like this, even though Harry really did get lost at Disneyland at the age of four, we wouldn't want to call that seeming memory a genuine memory. It lacks the reliably stable connection to the actual event.

- The second kind of case, taken from Martin and Deutscher's 1966 article, is described as follows:

> *[L]et us say that [Kent] has told his friend Gray what he saw of an accident in which he was involved. Kent has a second accident in which he gets a blow on the head which destroys all memory of a period in his past, including the time at which the first accident occurred. When Gray finds that Kent can no longer remember the first accident, he tells him those details which Kent had told Gray in the period between the first and second accidents. After a little while Kent forgets that anyone has told him about the first accident, but still remembers what he was told by Gray. It is clear that he does not remember the accident itself. ... Kent witnessed the first accident, can now recount what he saw of it, but does not remember it.*

- Martin and Deutscher introduce this example to demonstrate that the reliably stable connection component of the theory isn't enough. There is a reliably stable connection between Kent's experience of the first accident and his later memory of it, but the reliably stable connection consists in part in Gray's recollection and testimony to Kent as to what occurred in that first accident. What's missing from the case of Kent is the "continuous brain-based connection." This is why Kent's seeming to remember also can't count as a genuine instance of remembering.

- Because many philosophers found Reid's criticisms of Hume (and by extension Russell) compelling and because of the popularity of the causal theory of memory, in the years since the publication of Martin and Deutscher's article, a great many philosophers have come to think that memory has a very specific sort of limitation when it comes to providing support for knowledge.

- To see that, let's look back at some other sources of knowledge, such as self-awareness and sense perception.
  - If you're currently experiencing a certain mental state—say, a headache—and you ask yourself whether you are, you can form the belief that you are currently experiencing a headache. That's a new piece of information, one that you didn't previously have.
  - If you look outside and see that there are now two robins in the shrub outside your window, that's also a new piece of information. The only way that you could now know that there are currently two

robins in that shrub would be to check and see for yourself, or to have somebody else tell you.

- This suggests to philosophers that both self-awareness and sense perception are generative sources of knowledge—both can generate new knowledge, or give you knowledge that you didn't previously have.

- For a long time, the common view among philosophers has been that memory is not a generative source of knowledge. Instead, most philosophers hold that memory is a preservative source of knowledge. This means that you can't know something on the basis of memory unless you knew it at some previous time on the basis of some other source of knowledge.

- Take the case of your headache again. If your current knowledge that you had a headache at some earlier time is based on memory, then that can only be because at that earlier time you knew about your headache on the basis of self-awareness.

- Recently, however, some philosophers have begun to question the received view that memory is a purely preservative source of knowledge. For example, philosophy professor Jennifer Lackey has offered a number of examples of cases that she takes to illustrate the possibility of memory as a generative source of knowledge.

READINGS

Bernecker, *Memory*.

Locke, *Memory*.

QUIZ

1 Which type of memory do we rely on when we remember a specific fact, such as a date from history?

a Procedural

b Declarative

c Episodic

2 TRUE OR FALSE

The phenomena of aphantasia and severely deficient autobiographical memory provide an example of the independence of declarative and episodic memory.

3 The fact that H. M., who had no ability to store long-term semantic or episodic memories, was able to learn the mirror-drawing skill is an example of which of the following?

a The independence of procedural memory from the other types of memory

b The independence of episodic memory from the other types of memory

c The independence of semantic memory from the other types of memory

4 TRUE OR FALSE

Bertrand Russell thought that his five-minute hypothesis was evidence that memories are not dependent on the past.

5 TRUE OR FALSE

Thomas Reid objects that David Hume's view of memory—like Bertrand Russell's—is objectionable because it actually relies on the view of memory that it rejects.

6 Which view of memory suggests that it relies on a reliably stable, continuous brain-based connection to a past object or event?

a The Russell-Hume view

b Thomas Reid's view

c The Martin-Deutscher causal theory of memory

Answer key can be found on page 207.

LECTURE 10

Confabulations and False Memories

How does memory knowledge fit into the general picture of knowledge that has been developed so far? Is it best understood as having a coherentist or foundationalist structure? Can internalist or externalist theories explain how our memories provide us with evidence for knowledge?

Flashbulb memories involve the feeling of particularly clear memories that arise at a time of heightened public emotion—such as the JFK assassination, the moon landing, the O. J. Simpson verdict, or 9/11. But although these memories feel incredibly clear and precise, they're actually not. They can involve embellishments and alterations as they are recalled and re-recalled over time.

COHERENTISM AND MEMORY

- On the coherentist picture, what makes your true belief count as knowledge is that it forms part of a coherent framework of other beliefs that you have. In order to evaluate the coherentist picture of memory knowledge, there are two points to consider.
  - We need to distinguish a descriptive account of what goes on in remembering from a normative account of what would need to go on for that remembering to count as knowledge. Basically, a descriptive account says what actually happens while a normative account says what should happen. Applying that to memory, we need to distinguish two separate things: what we—or our brains—actually do in remembering a past object or event and what we would have to do to provide memory support for our knowledge.
  - In the lectures on how sense perception supports knowledge, we distinguished between tasks that you perform and tasks that only your brain performs. The distinction is that the tasks that only your brain performs, called subpersonal tasks, are ones that you're not aware of and that you have no control over.
- Researchers in the cognitive psychology and neurobiology of memory have a pretty clear idea

that the act of recalling a memory is less like pulling up an old file from a drawer and more like reconstructing a picture from stored pieces according to schematic instructions. In other words, the way our brains bring up memories of previous events fits well with the coherentist picture.

- In the 2010 article "Why Science Tells Us Not to Rely on Eyewitness Accounts" in *Scientific American Mind*, psychology professors Hal Arkowitz and Scott Lilienfeld review some of the literature on episodic memory, particularly as it relates to eyewitness testimony. They note:

> *The act of remembering, says eminent memory researcher and psychologist Elizabeth F. Loftus of the University of California, Irvine, is "more akin to putting puzzle pieces together than retrieving a video recording."*

- One issue with appealing to the way the brain puts puzzle pieces together to reconstruct memories is that this seems to be something that often just the brain does, rather than something that we do. In other words, it seems to be one of those cognitive processes that is subpersonal rather than personal. However, for a process to count as evidence for a coherentist structure of knowledge, the process would have to be a personal process, rather than a subpersonal one.

- The issue of the distinction between subpersonal and personal processes seems to pose a problem for the coherentist, but the next issue is, if anything, even worse. This involves the descriptive versus normative distinction.

- The problem is that what Loftus describes is how our brains actually reconstruct memories. However, it's a separate question as to whether that way of reconstructing memories is reliably accurate. And there is good reason to be skeptical about whether mere coherence is sufficient to guarantee accuracy.

- As Arkowitz and Lilienfeld emphasize, something as simple as questioning from a lawyer can change an eyewitness's memory of events. That's because the witness can unwittingly change the memory by adding information that the lawyer provides during questioning.

- The problem for the coherentist is that if you make having knowledge dependent on being able to spin a good yarn, then sometimes the best or most believable yarns aren't in fact true.

- An extreme example of what's wrong with the coherentist picture of memory knowledge is the case of confabulations—the stories that people unknowingly make up to fill in the holes in their memory. They are particularly common in people with cognitive or mental disorders, but otherwise healthy individuals also confabulate. It is important to remember that confabulation is a subpersonal process; the person isn't intentionally making things up and isn't aware that he or she is doing anything other than recalling a past event.

- It doesn't seem plausible to think of the structure of memory knowledge in coherentist terms.

FOUNDATIONALISM AND MEMORY

- There are two basic types of foundationalist theory of memory knowledge—the inferential theory and the naive foundational theory—and they are similar in structure to the two types of foundationalist theory of knowledge on the basis of sense perception.

- According to the inferential theory, the way that you acquire knowledge on the basis of memory is that you have a memory experience and then you argue that, given that experience, what you seem to remember probably did occur.

- For example, consider a case of discursive knowledge. Suppose you seem to remember that George Washington was the first president of the United States of America. You'd then reason as follows:

1. I seem to remember that Washington was the first president of the United States of America.
2. In the past, whenever I have seemed to remember something, it has usually turned out that what I remembered is true.
3. Therefore, it's probably true that Washington was the first president of the United States of America.

- This model seems very implausible as described. It just doesn't seem that memory works like that. But there's a more serious problem with this model: It is viciously circular. In order to establish the second step of the argument, you have to remember what happened in the past. This is the problem that Reid raised for the Hume and Russell model of memory presented in the previous

lecture. So, the inferential form of foundationalism is a nonstarter.

- According to the naive theory, memory gives you a good reason to believe something when you seem to remember it and you are not presently aware of anything that suggests that what you remember is false.

- In some ways, this seems a much more plausible picture of how memory actually works—as a personal process. In other words, when you ask yourself how it feels to remember something, it doesn't feel like you're putting puzzle pieces together, nor does it feel like you're formulating an argument. Instead, it feels like you remember something and—unless you have some specific reason to question that memory—it's okay simply to rely on what you remember.

- However, there are two very significant problems for the naive theory.
  - Consider a case in which you see some event but have reason at the time to doubt that what you're seeing is genuine. For example, suppose you see a very well-executed fake documentary in which a group of children seems to knock down a grown man and steal his briefcase, but you know that it's a fake documentary because you're viewing it in the context of a film festival of fictive documentary films. At a later date, you remember what you experienced previously, but you've forgotten the reasons you had at the time to doubt what you saw. As far as the naive theory is concerned, this would be a case in which you have good reason to believe the children really did knock down the grown man and steal his briefcase. That's because your memory now meets the two conditions required by the naive theory: You seem to remember the children committing the theft, and you are not presently aware of anything that suggests what you remember is false. This seems strange, however. It would be odd if your forgetfulness of the evidence rebutting your experience would actually help you. But that's exactly what the naive theory suggests in this case.
  - When we remember something, our brains fill in the remembered experience with plausible details. Rather than being a recording of the details of the experience that is simply recalled in the brain, a memory is more like a reconstruction. This is one of the ways that our memories can mislead us, and it is the source of one of the problems for the coherentist picture of memory.

Psychologists call our ability to remember where we acquired a particular piece of information source monitoring. The problem is that we're really bad at source monitoring. It's maybe not surprising that we're bad at remembering the specific source of information—say, whether we read something in *The New York Times* versus *The Wall Street Journal*. However, it turns out that we're even bad at remembering whether we actually read the information at all—rather than, say, hearing it during a 10-minute gossip session at the water cooler.

Even more dramatically, we often mistake the source of our memory as involving our own experiences when the source was actually something we read or heard or a photo we saw.

Our difficulties with source monitoring pose a problem for foundationalism. Part of the plausibility of naive foundationalism is that when you have a seeming memory, the memory experience you have isn't devoid of additional information. In other words, it's not merely a seeming memory experience of a past object or event (if it's an episodic memory) or of a supposed fact (if it's a declarative memory).

Instead, the seeming memory experience also contains information of its supposed source; that is, you remember it as something you experienced yourself or as something you heard or read. And the fact that you remember it as having come from a certain source seems to be part of what makes the supposed memory so believable.

But if it's true that we are subject to source monitoring failures we're unaware of, this suggests that the naive foundationalist theory is just that—naive. We should look for something better.

If what's been presented so far in this lecture is correct, neither form of internalism—neither coherentism nor foundationalism—does a good job of explaining how memory provides us with support for knowledge.

But there are some ways in which both coherentism and foundationalism add to our understanding of how memory contributes to knowledge.

It is in fact likely that the computational processes involved in the retrieval of memories from the information encoded in the brain operate according to coherentist principles. Because of this, it is plausible that an investigation of how coherence can contribute to accurate reconstructions of information will be useful for cognitive scientists. It will help them better understand how the brain supports the encoding, storage, and retrieval of memories. In other words, paying attention to the lessons we learn from coherentism can help us understand the subpersonal mechanisms involved in memory.

Furthermore, the way it feels to rely on memory for information seems to be best explained by a foundationalist account. For that reason, when it comes to explaining why someone might behave in a certain way on the basis of his or her memory, it would be good to pay attention to lessons we've learned from foundationalism. In other words, appealing to foundationalism can help us understand how memory contributes to the personal factors that can help explain someone's behavior.

The problem is that neither coherentism nor foundationalism is sufficient to explain how memory can support knowledge. That's because neither of them does a good job of accounting for if and when memories are reliably accurate. In order to do that, we'll need to turn to a consideration of externalism.

EXTERNALISM AND MEMORY

- Remember, there are two aspects that are essential to externalist theories.
  - They make reliable accuracy the central quality for determining whether a certain cognitive process is a source of knowledge. In the case of memory, this means that only those memory processes that reliably give us accurate information about the remembered facts, objects, or events will count as potential sources of knowledge.
  - They don't require that we be aware of the operations of the processes that give us information. This means that externalist theories—in contrast with internalist theories—allow that we can be unaware of the process we are employing to acquire information and ignorant of whether that process is in fact reliably accurate. In the case of memory, this means that we at least sometimes acquire information on the basis of memory processes we're unaware of or whose accuracy we cannot confirm.

- It seems that the phenomena of confabulation and source monitoring failures provide evidence for both of the aspects of memory that support externalism.
  - Both involve evidence that some cognitive systems are more reliably accurate in encoding, storing, or retrieving information in memory.
  - Both support the idea that we are often unaware of the processes that we actually employ in memory and that we are often equally unaware of whether or why the processes we employ are reliably accurate.

- This suggests that, as was the case with sense perception, externalism provides the best account for explaining the structure of knowledge in the case of memory.

- There's one further reason why the phenomenon of memory provides us with a strong motivation for embracing externalism. It's known within philosophy as the problem of forgotten evidence.

- Often, we retain a particular well-justified belief without remembering the evidence that originally led us to form the belief in the first place. Two related phenomena are forgotten negative evidence and source monitoring failures. But in each of these cases, the problem is that foundationalism is forced to say that a belief that shouldn't count as well justified does count as well justified.

- The problem of forgotten evidence is the mirror image of those problems for foundationalism. This is a

problem because the belief *is* well supported; it's just not well supported by any belief or experience that the believer still possesses. But because all internalist theories—whether coherentist or foundationalist—base the support for your belief on what's currently in your mind, they cannot account for a belief that is supported by evidence you no longer possess. Externalist theories, however, require a reliably accurate process to explain what makes a belief well supported. Because a process can involve components that are no longer present, externalist theories don't seem to have a problem with forgotten evidence.

All of this suggests that considerations of memory knowledge provide further support for externalist theories of knowledge in general.

READINGS

Bernecker and Michaelian, eds., *The Routledge Handbook of Philosophy of Memory*.

Michaelian, Debus, and Perrin, eds., *New Directions in the Philosophy of Memory*.

QUIZ

1 The fact that episodic memory works more like reconstructing a picture from stored pieces would seem to provide some support for which of the following?

a Coherentism

b Foundationalism

2 TRUE OR FALSE

The phenomenon of confabulation suggests that there are reasons for questioning whether coherentism can explain the reliability of memory.

3 TRUE OR FALSE

The problem for the inferential foundationalist memory theory is that it is viciously circular.

4 Which of the following would seem to pose a problem for the naive internalist foundationalist theory?

a The problem of forgotten evidence

b The problem of confabulation

c The problem of the unreliability of source monitoring

d All of the above

e None of the above

Answer key can be found on page 207.

LECTURE 11

The Extended Mind

With the rise of interest in the extended mind, philosophers and cognitive scientists have begun to question whether memory is limited to a brain-based connection to some past object or event or whether other types of processes for recording information outside of the brain could count as forms of memory. The 1998 paper in the philosophy journal *Analysis* written by professors of philosophy Andy Clark and David Chalmers entitled "The Extended Mind" gave the hypothesis its name. Even if you remain skeptical about the notion of extending the mind, the lessons you learn in this lecture will nevertheless be significant for the discussion of the structure of knowledge.

USE OF THE ENVIRONMENT AS A COGNITIVE AID

- Actions like using pen and paper for a mathematical calculation or using Scrabble tiles to aid in brainstorming possible words to spell are called epistemic actions by Clark and Chalmers. Epistemic actions are, as they put it, actions that "alter the world so as to aid and augment cognitive processes such as recognition and search."

- Some of the epistemic actions are such that they essentially involve environmental aids to successfully complete the cognitive task. Using a pen and paper for long division seems like one such action.

- Those cases involve what Clark and Chalmers call coupled systems, composed of a human reasoner and some portion of his or her environment. They define a coupled system as a case where

> *the human organism is linked with an external entity in a two-way interaction, creating a coupled system that can be seen as a cognitive system in its own right. All the components in the system play an active causal role, and they jointly govern behavior in the same sort of way that cognition usually does. If we remove the external component the system's behavioral competence will drop, just as it would if we removed part of its brain.*

- By appeal to the dual notions of epistemic actions and coupled systems, Clark and Chalmers define their position as active externalism. By this term, they mean that when we are engaged in epistemic actions, we (at least sometimes) count as coupled systems with external entities that are genuinely cognitive.

Clark and Chalmers sum up the central claim of the extended mind hypothesis like so: "If, as we confront some task, a part of the world functions as a process which, were it done in the head, we would have no hesitation in recognizing as part of the cognitive process, then that part of the world is (so we claim) part of the cognitive process."

- The main strategy Clark and Chalmers pursue in the extended mind hypothesis is argument by analogy. They try to present a case in which a process that extends outside of the brain has all the marks of a cognitive process and then argue that we should therefore recognize that case as a genuinely cognitive process.

- An analogy that Clark and Chalmers develop in some detail has two cases (although they also discuss a third case that we won't consider).
  - The first case involves Inga and a normal brain-based belief. Inga hears about an exhibition that she wants to see at the Museum of Modern Art. She recalls that MoMA is located on 53rd Street, so she heads there and waits in line to enter the museum. In other words, Inga represents a typical case of declarative memory.
  - The second case involves Otto, who has Alzheimer's disease. Like many people with Alzheimer's, Otto relies on environmental supports to provide him with aids in structuring his daily life. In particular, Otto carries a notebook with him that he uses as a support. Whenever he learns something new and significant, he records it in the notebook. If he needs some information, he looks it up. Otto hears about the exhibition at the MoMA and also wants to see it. He looks up the address of MoMA in his notebook and reads in the notebook that it is located on 53rd Street. So, Otto heads to 53rd Street and waits in line to enter the museum.

- What Clark and Chalmers want to underscore is that Inga's brain-based belief and Otto's notebook-based belief function the same way in explaining their actions. Just like Inga, Otto headed toward 53rd Street because he wanted to go to the MoMA and that's where he believed the MoMA to be. In Inga's case, philosophers would naturally say that she possesses the dispositional belief that the museum is on 53rd Street even when she's not considering it. So, Clark and Chalmers argue, we should also say that Otto believes that the museum is on 53rd Street even when he isn't consulting what he's recorded in his notebook. The key point, according to Clark and Chalmers, is that the information that Inga and Otto have stored functions in exactly the same way for both of them. For that reason, it shouldn't matter if that information is stored in the brain or in a notebook.

Clark and Chalmers suggest that there are no relevant differences between Inga's brain-based memory belief and Otto's notebook-based belief. Otto's external "memory" is just as integral to his daily functioning as Inga's internal, brain-based memory is to hers.

- But are there really no relevant differences between the Inga case and the Otto case?

- Clark and Chalmers suggest that we might raise four worries about whether the cases really are that similar. There might be differences in reliability, stability, accessibility, or phenomenology.

- The first two potential differences between Inga and Otto are reliability and stability. You might worry that Otto's method for remembering using his notebook is less reliable or stable than Inga's. This, however, seems to be obviously wrongheaded. If anything, the method of using a notebook as an external "memory" is more reliable and stable than using your internal, brain-based memory—not less.

- Of course, using the notebook as an external memory does open you to the possibility that it might get lost. But this possibility exists for your brain-based memory as well. Not only do we often lose individual memories, but there is also the more dramatic possibility that, through disease or injury, we could experience more significant memory loss.

- The other two differences—phenomenology and accessibility—are more serious.

- In the case of phenomenology, it doesn't seem that there is a relevant difference between the phenomenology of brain-based declarative memory and Otto's notebook-based memory. The reason for this is that there is no phenomenology of brain-based declarative memory.

- Suppose someone asks you who the 13th president of the United States was. When you remember a fact or piece of information, it just comes to you—or it doesn't. Perhaps you have the feeling of making an effort to remember, but there's no special feeling associated with remembering the information itself.

- There does seem to be a difference in phenomenology for episodic memory. Brain-based memories in which you imaginatively "relive" a past experience seem to be more difficult to "offload" onto the environment. But perhaps with the rapid improvements of interactive technologies, even that difference will become less significant.

- The difference in accessibility is a more serious objection. Inga can access her memory in the dark or when her hands are full, whereas Otto cannot access his notebook in those types of situations.

- Clark and Chalmers recognize the seriousness of the accessibility objection. They grant that the brain and body can be thought of as a self-contained set of fundamental cognitive resources. Furthermore, we obviously take our brains and bodies with us wherever we go, so there's no question that those resources are optimally accessible to us. Given this fact, it seems natural to distinguish cognitive tasks that we can perform with our brains alone—or even with the aid of our fingers, say, in the case of arithmetic—from cognitive tasks that we perform with a portable electronic device.

- Nevertheless, Clark and Chalmers don't think that the accessibility issue ultimately is a problem for the extended mind hypothesis.

EXTERNALISM AND THE EXTENDED MIND

- Recent studies seem to demonstrate that there is growing empirical evidence that fits well with the extended mind hypothesis. But even if you are still unsure about whether you find the hypothesis ultimately convincing, the question of whether the extended mind or extended cognition is a genuine phenomenon is interesting in its own right.

- Certainly, it might be very relevant to the specific topic of memory knowledge that we've been examining over the past few lectures. But the question is not all that central to the more general topics we've been discussing. Instead, we can take lessons from our discussion of the phenomena surrounding the topic of the extended mind for the more general questions concerning knowledge, regardless of how we feel about the specific issue of whether the mind really is—or can be—extended.

- If our discussion of the extended mind has taught us anything, it's that the specific process that leads to forming a belief matters. Suppose you're trying to evaluate whether to believe your friend Bill when he tells

you that 53,735 divided by 29 equals 1,852 with a remainder of 27. Unless you know that your friend is some sort of savant, it will make a difference to you if you know whether he solved the problem using pencil and paper or had to try to do it in his head.

- In the case of long division, in other words, the process of using pencil and paper makes you more reliably accurate in your calculations. For that reason, an externalist about knowledge will want to describe the process you employ when you solve a long division problem using pencil and paper as a different process than the one that you employ if you're forced to try to solve the problem in your head. And whether your belief about the solution to a particular long division problem counts as knowledge will quite likely depend on which of those processes you use.

- This means that the externalist has an advantage when describing the kinds of knowledge you acquire when you employ processes that extend out into the world outside of your brain and body. And because the claim that the process extends out into the world doesn't depend on any claims about the extended mind or extended cognition, the externalist has this advantage even if you insist that the mind is not extended but is limited to the confines of your brain—or your brain and body.

READINGS

Adams and Aizawa, *The Bounds of Cognition.*

Menary, ed., *The Extended Mind.*

QUIZ

1 Which of the following is an example of using the environment as a cognitive aid?

a Using a pencil and paper for long division

b Rearranging Scrabble tiles to see what words you can spell

c Using a nautical slide rule to compute problems involving time, speed, and distance on the water

d All of the above

e None of the above

2 TRUE OR FALSE

Andy Clark and David Chalmers call systems in which human reasoners are linked with some portion of their environment in a two-way interaction for the purposes of information processing a coupled system.

3 TRUE OR FALSE

Despite the similarities between Inga's normal, brain-based belief and Otto's atypical, notebook-based belief, Andy Clark and David Chalmers ultimately think that Otto's belief is too different from Inga's to count as a genuine case of the extended mind.

4 Which of the following are potential worries for the possibility of the extended mind?

a Reliability and stability

b Accessibility

c Phenomenology

d All of the above

e None of the above

5 TRUE OR FALSE

The results of Betsy Sparrow, Jenny Liu, and Daniel Wegner suggest that Andy Clark and David Chalmers's thesis of the extended mind is similar to the psychological phenomenon of transactive memory.

Answer key can be found on page 207.

LECTURE 12

Do We Have Innate Knowledge?

Consider the claim that equilateral triangles have three equivalent angles. Presumably, you can figure out that this claim is true. But how? Certainly, you don't have to draw a triangle and examine it to check. And even if you did draw a triangle to check, that would actually only tell you that at least one equilateral triangle—the triangle you just drew—has three equivalent angles. It wouldn't give you evidence that all equilateral triangles have three equivalent angles. It turns out that there are deep philosophical issues surrounding the question of how we know the truth of the claim that all equilateral triangles have three equivalent angles and other claims like it.

RATIONALISTS VERSUS EMPIRICISTS

- The traditional philosophical controversy about how we know that all equilateral triangles have three equivalent angles involves a debate between two diametrically opposed schools of thought. One school of thought, the rationalists, argue that humans possess a distinct faculty of reason and that humans could know truths of reason—such as the truth about the angles of an equilateral triangle—by exercising their faculty of reason.

- In fact, at the extreme, rationalists believe that all knowledge is derived from reason. For example, Descartes holds that in order to have knowledge of something, you have to have infallible evidence—evidence that could not be false. Because the senses are fallible, Descartes does not think that the senses can provide you with knowledge. Only reason can provide you with infallible evidence. So, for Descartes, only reason can provide you with knowledge.

- Descartes and the other rationalists seem to do okay with our knowledge of the truths of reason, but they don't do so well with our knowledge of facts about the world, such as whether it's hot outside. These facts do not seem like the types of things you could know by reasoning your way to an answer. You would have to look and see how the world actually is—or remember what you saw when you did look.

- On the opposite end of the spectrum from the rationalists are the empiricists, including such philosophers as John Locke and David Hume—according to whom all knowledge is derived from experience.

- Their situation looks like a mirror image of the rationalists' situation. The empiricists seem to do well with our knowledge of facts about the world. It seems pretty obvious to think that those are the kinds of things that we learn from experience. But you can know about the angles of an equilateral triangle without ever seeing a triangle.

- This is how the debate stood until the middle of the 18th century, when Immanuel Kant entered the picture. He suggested that the debate between rationalism and empiricism was a fruitless one because the two sides had thought about their disagreement too simplistically. Kant hoped to show that by thinking more clearly about what the debate involved, we would also be able to find a solution to the conflict.

Defenders of rationalism include René Descartes, Benedict de Spinoza, and Gottfried Wilhelm Leibniz; defenders of empiricism include Thomas Hobbes, John Locke, George Berkeley, and David Hume.

KANT'S DISTINCTIONS

- Where rationalists thought in terms of truths known by virtue of reason alone and empiricists thought in terms of truths known in virtue of experience alone, Kant suggested that we introduce two different distinctions: one between analytic and synthetic claims and another between a priori and a posteriori knowledge.

- The distinction between analytic and synthetic claims has to do with whether the truth of a claim depends on the meaning of the words used to make the claim.

- An analytic claim is a claim that you can know to be true solely on the basis of the meaning of the words in it. For example, if you know what "poodle" means and you know what "dog" means, then you know that the claim "all poodles are dogs" is true.

- Contrast that with the claim "some poodles are white." That's a synthetic claim. It's also true, but you can't know that it's true simply by knowing the meaning of the words in the claim. You also have to actually get out in the world and look at some poodles.

- The second distinction is one that gets directly at the disagreement between rationalists and empiricists. It's based on whether the evidence you have comes from experience.

- The knowledge that we have of things on the basis of specific experiences is called a posteriori knowledge by Kant. If you look outside and see that it's sunny outside where you are, that's a posteriori knowledge. As a nod to the empiricists, Kant was happy to admit that we have that kind of knowledge.

- However, Kant also thought it obvious that some things you know are based on evidence that doesn't come from experience. He called those cases of a priori knowledge—knowledge that you have independent of any specific experience that supports it.

- Empiricists can be happy with some of that type of knowledge. Take the claim "all poodles are dogs." Suppose you grew up on an island where there were no poodles, let alone any other type of dog. You've never seen a canine of any type. Nevertheless, if you know English, you can still know that all poodles are dogs.

- In other words, some of our a priori knowledge is knowledge that we gain simply by knowing the meaning of words, even if we don't have any specific experience with the things or qualities named by those words.

- This sort of a priori knowledge also involves analytic claims—claims that you can know are true simply because of the words they involve.

- This might seem a bit strange, but it shouldn't be. There are many cases in which we can learn a lot about something without directly experiencing it, simply by learning facts about language.

- Kant thought that there was knowledge that we have independently of experience that isn't simply implied by the meanings of the words we use to state that knowledge. In particular, Kant thought that there were two big branches of knowledge that involved the synthetic a priori: mathematics and philosophy.
  - According to Kant, when you discover that 17,361 divided by 27 equals 643, you're doing something more than simply relying on the meanings of the terms in that statement. Another way of saying this is that you can understand the meanings of 17,361, 27, and 643 and still wonder whether 17,361 divided by 27 is equal to 643. That's not the case with analytic truths, according to Kant. If you understand the meaning of "poodle" and "dog," you won't still wonder whether all poodles are dogs.
  - Kant believed that if he could convince you that mathematics involved synthetic a priori claims, then you would have no reason to argue against the idea that philosophy involves synthetic a priori claims as well. In other words, Kant saw his argument for the existence of the synthetic a priori as providing proof of philosophy as a well-founded discipline, on a par with mathematics.

A PRIORI SYNTHETIC TRUTHS

- Suppose you agree with Kant that we do have at least some a priori knowledge. Where does that knowledge come from? How do we know, for example, that all equilateral triangles are equiangular—particularly because we can know that without physically examining any equilateral triangles?

- The point of introducing the synthetic a priori is to reject one way of answering that question. For the empiricist, one thing that would be acceptable would be to say that you can have a priori knowledge as long as it is knowledge derived from the meaning of words or concepts. So, as long as all a priori knowledge is analytic, then it would be okay for the empiricist.

- Once you introduce the synthetic a priori, however, it is no longer as easy for the empiricist to explain where that knowledge comes from—the synthetic a priori knowledge that doesn't ultimately come from knowledge of meaning.

- So, where does this knowledge come from? Let's examine three answers.

- The first suggestion for where our a priori knowledge comes from is a form of rationalism. It's less extreme than Descartes's rationalism, because it doesn't claim that we learn everything through reason. Instead, these more modest rationalists just suggest that it is reason that explains how we learn the a priori synthetic truths of mathematics and philosophy. These modest rationalists suggest that we recognize those truths through the use of a rational faculty, or faculty of reason. There are two big problems with this idea.
  - The first problem has to do with how the faculty of reason comes to know the things it knows. Modest rationalists like to suggest that reason "recognizes" or "perceives" a priori synthetic truths. But this is just a metaphor, and it isn't helpful. Modest rationalists explicitly reject the notion that a brain-based explanation of how we come to know a priori synthetic truths is something we ought to seek. That's just not a strategy that is worth pursuing.
  - The second problem is that modest rationalists are very unclear about what a faculty of reason might be. The one thing that modest rationalists do seem to suggest is that the faculty of reason can't just be a mechanism or collection of mechanisms in the brain. And rejecting an appeal to brain-based mechanisms for understanding human knowledge isn't a promising explanatory. But even if you're

not opposed in principle to explanations that ignore the role of the brain and brain mechanisms in knowledge, certainly you shouldn't be satisfied with a name that stands for nothing. And unfortunately, in the case of modest rationalism, it seems that that's all we have. Just a label—the faculty of reason—with no further detail about what such a thing might be.

- Modest rationalism doesn't seem very plausible. Luckily, the other two explanatory strategies—the innate strategy and the language strategy—are based on mechanisms in the brain. Furthermore, the strategies are compatible with each other, so we can appeal to either one or both of the strategies to help explain our knowledge of a priori synthetic truths.
  - The innate strategy involves innate brain mechanisms. Such mechanisms almost certainly play a role in our knowledge of basic facts about numbers.
  - According to the language strategy, mathematical thinking is a result of our use of language.
- Both of these suggestions are plausible. The good news is that we don't need to choose between them; we can accept that both innate brain structures and language play a role in synthetic a priori knowledge.

The evidence that innate mechanisms in the brain play a role in our knowledge of mathematics is strong. Beginning in the 1980s, developmental psychologists introduced what became known as the violation-of-expectation paradigm to investigate the innate cognitive abilities of preverbal children, even very young infants. Using that research technique, researchers have been able to establish that very young infants—as young as three or four days old—have a number of innate mental abilities that are relevant to our discussion of a priori synthetic knowledge.

Newborns have the ability to recognize when different collections of objects have differing numbers of objects. Psychologists refer to this as the ability to subitize, and developmental psychologists suggest that this ability is evidence that infants have an innate conception of number. In fact, by the age of a few months, babies behave as if they have an intuitive understanding of very simple arithmetic operations, such as 1 + 1 = 2.

One of the pioneers of the suggestion that our ability to think abstractly is a result of our use of language was Soviet psychologist Alexander Luria. In a groundbreaking series of studies, he interviewed a number of Central Asian peasants and found that the ability to think and reason abstractly depended on the level of language learning that the peasants had attained.

Research conducted by cognitive neuroscientist Stanislas Dehaene shows that there are at least two separate brain systems responsible for mathematical cognition: one involves approximate arithmetic and the other is employed when we compute exact arithmetic. Dehaene suggests that this second system is tied to the language that a person has learned.

This indicates that the best explanation of our a priori synthetic knowledge—such as our knowledge of mathematics—might involve both innate brain structures that provide us with the most fundamental concepts that we depend on for such knowledge and structures contained within our language that allow us to build on those fundamental concepts.

READINGS

Casullo, *A Priori Justification*.

Moser, ed., *A Priori Knowledge*.

QUIZ

1 The fact that Nicholas Saunderson, the famous blind Lucasian professor of mathematics at Cambridge, could know facts about geometry would seem to provide some evidence against which of the following?

a Pure rationalism
b Pure empiricism

2 Which of the following did Immanuel Kant think involved claims that are true in virtue of logic and meaning alone?

a A priori
b A posteriori
c Analytic
d Synthetic

3 Which of the following did Immanuel Kant think involved claims that were known independently of experience?

a A priori
b A posteriori
c Analytic
d Synthetic

4 According to Kant, philosophical or mathematical truths are some of the few claims that can both be which of the following?

a Analytic a priori
b Synthetic a priori
c Synthetic a posteriori

5 Karen Wynn's experiments on the mathematical abilities of very young infants provide at least some evidence that some human mathematical knowledge is which of the following?

a The result of rational intuition
b Due to innate mathematical brain modules
c Due to linguistic knowledge

6 Soviet psychologist A. R. Luria's research suggests that at least some human mathematical knowledge is which of the following?

a The result of rational intuition
b Due to innate mathematical brain modules
c Due to linguistic knowledge

Answer key can be found on page 207.

LECTURE 13

How Deduction Contributes to Knowledge

Much of what we know is not something that we've experienced directly. In many cases, what we know is inferred. In fact, as in the case of the foundationalist position, many philosophers think that inference is essential to the structure of knowledge.

DEDUCTION VERSUS INDUCTION

- Deduction involves inference based on the purely logical relationships that hold among the steps that make up the inference. A classic example is the syllogism, a form of inference that Aristotle studied more than 2,000 years ago. This syllogism is a paradigmatic example of deductive inference:
  1. All men are mortal.
  2. Socrates is a man.
  3. Therefore, Socrates is mortal.

- Consider this significant attribute of the inference. There is no way for the steps of the inference—what philosophers call the premises—to be true and the conclusion to be false. The conclusion, in other words, necessarily follows from the premises.

- This is the central attribute of deduction. A well-formed deductive inference is one in which the conclusion necessarily follows from the premises. Such well-formed inferences are referred to as deductively valid.

- This is a remarkable property of deductive inference. It offers a guarantee that you can be sure, purely based on the form of the inference, that if you have true premises, you'll have a true conclusion.

- This remarkable power, however, comes with two significant costs.
  - Deductive validity gives us a guarantee that if our premises are true, our conclusion will also be true, but if any of the premises are false, then all bets are off—the conclusion could be false, but it could also be true.
  - The conclusion of a deductively valid argument will always contain no more information than the premises of the argument. In fact, the conclusion of a deductive argument generally contains dramatically less information than the premises of that argument. You could actually think of this as the function of a deductive argument: Its conclusion allows you to focus on the particular implication of the information contained in the premises that interests you.

- If what you want is an argument whose conclusion contains information that is not contained in the premises, then what you want is a nondeductive argument. There are a number of different types of nondeductive arguments, but we'll focus on inductive arguments.

- Inductive arguments are ones whose premises merely make their conclusion highly probable. Consider the following argument.

1. In the past, I have often observed cardinals, and they have always been red.
2. Therefore, the next time I observe a cardinal, it will also be red.

- You can think of the first step of the argument as a shorthand for a list of all of the individual occasions on which you've observed cardinals and they've been red. If you think of that step in this way, then the support for the conclusion of the inductive argument—that the next cardinal you see will also be red—will presumably be stronger the more observations of cardinals you've had in the past.

- This fact about inductive arguments allows us to appreciate the key aspect about them that distinguishes them from deductive arguments: Inductive arguments can involve weaker or stronger support for their conclusion. The strength of inductive arguments depends on the level of support that their premises provide for the conclusion, whereas deductively valid arguments with true premises always guarantee the truth of their conclusion.

- The trade-off for the fact that inductive arguments cannot guarantee their conclusions is that inductive arguments can involve conclusions that provide information not contained in the premises. In fact, any generalization in science—for example, that all electrons have less mass than any proton—is the conclusion of an inductive argument. This is because any general statement would necessarily go beyond the individual observations on which it must be based.

DEDUCTIVE ARGUMENTS

- Deductive arguments are central to both foundationalist and coherentist internalist theories.
  - In the case of Descartes's form of internalist foundationalism, deduction is the only way in which knowledge can be supported by the foundational beliefs you have. This is because Descartes required that knowledge be infallible, and only deductive inference provides a guarantee that its conclusions are true. But even other forms of fallible foundationalism still regard deductive inference as an important source of knowledge.
  - It's not exactly clear what goes into the relation of coherence. However, one of the components would have to involve something like logical consistency, and logical consistency involves deductive inference.

In other words, particularly for the internalist, deductive inference is of central importance. It is surprising, then, that philosophers have often been blind to a number of complications having to do with deductive inference. We'll focus on three representative problems.

The logical connectives used in deductive reasoning are defined in terms of how they affect the truth values of sentences.

If you add the logical "not" to a true sentence, you get a false sentence. And if you add the logical "not" to a false sentence, the result is true.

The logical "and" joins two sentences. The resulting compound sentence is true only when both sentences joined by the "and" are true. In all other cases, the "and" is false.

The logical "or" is also used to join two sentences, but in the case of the "or," the resulting sentence is true as long as at least one of the components is true. The sentence is also true when both are true. This is because the logical "or" is called the nonexclusive "or." That means if a logician wants to use "or" to mean "one or the other, but not both," then he or she explicitly has to stipulate that.

The "if-then," also called a conditional statement, also joins two sentences—like the logical "and" and the logical "or." Unlike these two, however, the position of the sentences joined by the "if-then" matters. Let's call the "if" sentence the antecedent (A) and the "then" sentence the consequent (C). If A, then C is only false when the antecedent A is true and the consequent C is false. In all other cases, the "if-then" is true.

- Despite the fact that some of these difficulties have been studied since the 1960s, psychologists and philosophers still have differing explanations for why these problems in deductive reasoning persist. However, there is universal agreement that the existence of these difficulties provides proof that humans are not naturally gifted at abstract deductive reasoning. Even highly intelligent people make very basic mistakes when they attempt to think deductively.

- An example demonstrating the fact that humans are not innately strong abstract reasoners is the Wason selection task, a psychological experiment first conducted by psychologist Peter Wason in the mid-1960s. Here's a description from an article cowritten by Wason and his collaborator Philip Johnson-Laird:

> *You are presented with four cards showing, respectively, "A," "D," "4," "7," and you know from previous experience that every card, of which these are a subset, has a letter on one side and a number on the other side. You are then given this rule about the four cards in front of you: If a card has a vowel on one side, then it has an even number on the other side. Next you are told: Your task is to say which of the cards you need to turn over in order to find out whether the rule is true or false."*

- As Johnson-Laird and Wason emphasize,

> *Very few highly intelligent [people] get the answer right spontaneously; some take a considerable time to grasp it; a small minority even dispute its correctness, or at least remain puzzled by it.*

- In their words, this is the answer to the selection task:

> *The most frequent answers are "A and 4" and "only A." They are both wrong. The right answer is "A and 7" because if these two stimuli were to occur on the same card, then the rule would be false but otherwise it would be true.*

- The results of the Wason selection task are some of the most widely replicated in the history of modern psychology. When presented with abstract problems testing their intuitive grasp of "if-then" deductive reasoning, only less than 25 percent—in some experiments as few as 7 percent—of people complete the selection task correctly.

- Later experiments have found that the subject matter can make an enormous difference in how well people do in completing the selection

task. The issue of why and how different versions of the selection task can yield such different results still has researchers puzzled.

- But what is unquestioned is that humans are abject failures, by and large, at the abstract version of the selection task. And given that for formal logic the content of the rules shouldn't matter—only their structure should matter—this shows that humans are not naturally gifted at deductive reasoning.

- According to further research conducted by Wason, it turns out that humans are not all that good at applying the rules governing "or" and "and" statements either.

- If we're not all that good at deductive logical reasoning, why would traditional internalists—both foundationalists and coherentists—make deductive logical reasoning an essential component of justification? If the goal of epistemology is to explain how we have knowledge, rather than to demonstrate that we lack knowledge, then requiring deductive logical ability in order to have knowledge seems like a bad strategy.

- When most people think about themselves, they likely tend to think that they're actually pretty good at logical reasoning. In fact, one of the surprising aspects of Wason's research is that they contradict our self-image as strong logical reasoners.

- If many of us are in fact so poor at reasoning deductively, why are we not more aware of that fact? The answer can be found in one of the most famous social psychological experiments of the last quarter century.

- The experiment was performed by a professor of psychology at Cornell University, David Dunning, and a graduate student in psychology, Justin Kruger. Their article reporting on the experiment, "Unskilled and Unaware of It," which appeared in 1999 in the *Journal of Personality and Social Psychology*, has racked up more than 4,000 citations on Google Scholar.

- In the version of their experiment that is relevant to deductive reasoning, Dunning and Kruger gave a group of Cornell undergraduate psychology students a 20-item logical reasoning test drawn from study questions for the LSAT.

- On average, the students placed themselves in the 66th percentile on the logical reasoning test. Remember that the average has to be the 50th percentile.

- Dunning and Kruger tried to figure out the source of this misdiagnosis of logical ability, so they divided the test takers into quartiles and found

that the top 25 percent of the students actually underestimated their performance on the test.

- What accounted for the overestimate of logical ability on the part of the student participants was the worst performers. In particular, the bottom 25 percent of the students, who scored on average in the 12th percentile, believed themselves to be in the 68th percentile of all test takers.
- Dunning and Kruger suggest that "those with limited knowledge in a domain suffer a dual burden: Not only do they reach mistaken conclusions and make regrettable errors, but their incompetence robs them of the ability to realize it." This phenomenon has become so widely recognized that it has entered the literature as the Dunning-Kruger effect.

Both internalist foundationalists and coherentists require deductive reasoning as a central feature of how we are supposed to construct our knowledge. However, we are not only bad at deductive reasoning, but also insensitive to what poor logical reasoners we are. This is a bad combination for internalist theorists.

According to the externalist, there is no requirement that you can know something only if you can support your knowledge on the basis of reasons you have access to. That requirement—coupled with the suggestion that the way you must support your knowledge will at least sometimes involve deductive inference—is what leads to problems for the internalist.

Without that requirement, the externalist faces no challenge due to the fact that our deductive reasoning skills are not as good as we intuitively assumed. In this way, the externalist theorist is able to have a more clearheaded assessment of our actual reasoning abilities.

A consideration of deductive reasoning suggests at least some support for an externalist theory.

READINGS

Harman, *Change in View*.

Mercier and Sperber, *The Enigma of Reason*.

QUIZ

1 Inferences that have the property of having true conclusions whenever all of their premises are true are which of the following?

a Deductively valid

b Inductively strong

2 Generalizations in science are examples of the conclusions of which type of arguments?

a Deductive

b Inductive

3 Suppose it's true that it's sunny in Seattle and it's false that it's raining in Tucson. Which of the following sentences is true?

a It's sunny in Seattle, but it's not raining in Tucson.

b Either it's raining in Seattle or it's not raining in Tucson.

c If it's raining in Tucson, then it's cloudy in Seattle.

d All of the above.

e None of the above.

4 Suppose you are told to evaluate the rule that if a basketball player is a forward, then the player is more than 6 feet tall. Which of the following would you need to examine to know if the rule is true?

a The height of a guard

b The position of a player who is 5 feet 9 inches tall

c The position of a player who is 6 feet 2 inches tall

d All of the above

e None of the above

5 Which of the following is an example of the Dunning-Kruger effect?

a People who are the worst at logical inference think that they are among the best at such inferences.

b People are not good at understanding conditional statements.

c Valid deductive arguments cannot give you more information than you put into them.

Answer key can be found on page 207.

LECTURE 14

Hume's Attack on Induction

There is a long philosophical tradition of questioning whether inductive inference is at all useful in generating knowledge. We can also ask the separate question of whether humans are actually good inductive reasoners. This means that we have two challenges to consider: whether inductive inferences are ones that we should rely on in increasing our knowledge and whether we humans are in fact able to reason inductively in ways that increase our knowledge. This lecture focuses on the first question while the next tackles the second.

INDUCTIVE ARGUMENTS

- Recall that two of the main features of deduction are that it is certain and it can't add to the information that we already have. Induction differs from deduction along both of these dimensions.

- First, induction is not certain. Inductive arguments involve probability. The point of introducing the notion of probability is to be able to discuss events that may not be certain—events that have a chance of occurring but also a chance of failing to occur.

- Suppose you have a friend who smokes a pack of cigarettes a day. You tell him that if he keeps that up, he's likely to develop serious health problems and to die prematurely. He responds by saying, "My grandfather smoked a pack of cigarettes a day, never had any health issues, and lived to be 100 years old."

- An inductive argument—even a good one—doesn't guarantee that its conclusion will apply in every case.

- The argument you give to your friend would look like this:
  1. Many of those people who smoked a pack of cigarettes a day or more developed serious health problems and died prematurely.
  2. Therefore, if you smoke a pack of cigarettes a day or more, you will likely develop serious health problems and die prematurely.

- Much of the strength of this argument will depend on the support you have for the first claim. And that support will require at least two things.
  - Presumably, the larger the percentage of people who develop health problems when smoking, the stronger the link is between the data you've gathered about smoking in general and the chance that your friend, in particular, will also suffer negative health effects from smoking.
  - The more evidence you've gathered that links smoking to poor health, the better the support for your first claim will be.

- Suppose you're satisfied that you have good evidence for your first claim. Unlike in the case of a deductive argument, even an inductive argument that seems strong doesn't guarantee the truth of the conclusion given the truth of the premises.

- At best, your first claim just tells you that you have a great deal of evidence that there has been a strong connection observed between smoking and poor health. This observed connection, however,

cannot guarantee that there actually is a connection between smoking and ill health, nor can it guarantee, for example, that smoking causes ill health. It merely makes it likely that such a connection is real and that such a causal relationship exists.

- This is the first complication: The truth of the premises of an inductive argument do not guarantee the truth of the argument's conclusion, even when the argument enjoys a great deal of evidential support.

- Here's the second complication. In almost all inductive argument conclusions, there will be an element of probability in the statement of the conclusion. For example, you argued that your friend will *likely* develop serious health complications and die prematurely. The problem here is that your inductive evidence could be strong and the conclusion could even be true, but your friend still might not develop serious health complications and die prematurely.

- Sometimes the less likely event does in fact occur. This means that it is much harder to demonstrate that the conclusion of an inductive argument is false. Suppose you are told only that a certain event is likely to happen—or even highly likely to happen. The mere fact that it doesn't happen on a single occasion doesn't actually demonstrate that the claim was false. It's only by looking at vast amounts of data that you could assess whether the estimate of the likelihood was accurate or not. And if you're dealing with a one-off event, that's not always terribly helpful.

> Inductive arguments are much less straightforward and much more difficult to deal with than deductive arguments.

- Given the problems with inductive arguments, you might suggest that we simply stop dealing with them. In some ways, this is a typical philosophical reaction. When you read most works of philosophy, all you'll encounter are deductive arguments, not inductive ones. The reason for this is at least in part that philosophers prefer the clean simplicity and elegance of deductive arguments to the messiness introduced by inductive arguments.

- The reason inductive arguments are useful, however, is that—unlike deduction—inductive arguments can actually add new information to the information that we already have.

- Suppose you notice, on a number of occasions, that whenever you light a campfire, smoke rises from the logs. You find yourself hiking in the woods during a forest fire alert and see, over the rise of the next hill, smoke wafting up. You think, *Where there's smoke, there's fire*, and head back to your car to alert the park rangers. That thought goes beyond any experience you've ever had. All of the experiences you've had have been of particular events, including particular campfires, so all of your evidence has been of the form of specific camping trips and campfires.

- So, that general claim—that where there's smoke, there's fire—is new information. It's not actually information you've ever simply perceived, although it is based on information that you've perceived—namely the information about your previous experiences with fire and smoke.

- That's the power of induction. It can allow you to derive new information from your existing storehouse of experience. And then, using that new information you've derived, induction allows you potentially to predict future observations or to manipulate future events.

HUME'S INDUCTIVE CHALLENGE

- Given the great power and usefulness of inductive arguments, it's no wonder philosophers have had a great interest in showing that those arguments work and how they can be improved.

- One of the great challenges faced by those who attempt to explain and defend the use of inductive inference was given by David Hume in a series of works. In the first of those works, *A Treatise of Human Nature*, he introduced an attack on induction that became one of the most influential arguments in the history of philosophy.

- The goal of Hume's argument is to show that, despite its power and usefulness, inductive arguments are not in fact based on reason at all. According to Hume, the correct conclusion to draw about induction is that it provides us with no reason for our beliefs or actions at all!

- How does Hume establish such a seemingly radical conclusion?

- He begins by noting that in order for inductive arguments to be successful, the uniformity principle—which says that unobserved instances resemble

observed instances—must be true. The uniformity principle requires this: For an inductive argument to be successful, past performance must at least resemble the future results in ways that make generalizations about those future results possible.

- Hume's uniformity principle seems to be extremely plausible. Inductive reasoning can only work as long as the unobserved cases that you're trying to predict resemble the observed cases on which you base your argument.

The uniformity principle is one of the basic assumptions at the root of scientific inquiry. As astrophysicist Brian Greene noted,

> *Science is not describing a universe out there, and we're separate entities .... We're part of that universe, we're made of the same stuff as that universe, of ingredients that behave according to the same laws as they do elsewhere in the universe.*

- Hume's argument is as follows:
  1. In order for someone to have reason to believe any matter-of-fact claim about the unobserved, he or she must first have reason to believe that the uniformity principle is true. To make this claim plausible, it would help to think from an internalist perspective. The uniformity principle is neither self-evident nor something we can know on the basis of immediate experience. Therefore, if you're an internalist—whether a foundationalist or a coherentist—Hume's first claim seems very plausible.

If you're an internalist, then you think that everything you have reason to believe has to be supported by something else you also have reason to believe, with a few potential exceptions.

If you're a coherentist, then there are no exceptions; everything you have reason to believe has to be supported by its coherence with everything else you have reason to believe.

If you're a foundationalist, then the exceptions are the foundational beliefs. Typically, for foundationalists, this would include things that are self-evident or things you believe on the basis of immediate experience.

2 All claims are either what Hume calls relations-of-ideas claims or matter-of-fact claims. Matter-of-fact claims can be either true or false. Relations-of-ideas claims are made true by the concepts, or ideas, they involve; they are claims you can know without going out into the world and checking. The claim that all statements can be divided into one of these types, sometimes called Hume's fork, seems at least somewhat plausible.

3 The uniformity principle is not a relations-of-ideas claim. There is nothing about the meaning of the ideas involved in the claim that unobserved instances resemble observed instances that would guarantee the truth of that claim. If all claims are either relations-of-ideas claims or matter-of-fact claims and if the uniformity principle is not a relations-of-ideas claim, then the uniformity principle must be a matter-of-fact claim.

4 Because the uniformity principle says that unobserved instances resemble observed instances, that means the uniformity principle is a claim about the unobserved.

5 Because the uniformity principle is itself a matter-of-fact claim about the unobserved, that means that in order for you to have reason

to believe that the uniformity principle is true, you must already have reason to believe that the uniformity principle is true.

- Because the support for the uniformity principle comes from the uniformity principle itself, the justification for the uniformity principle is viciously circular. And if that's the case, then Hume has the conclusion that he's after: We have no good, noncircular reason for believing in the conclusions of inductive arguments.

> Because pretty much all of science depends on inductive reasoning, in a nutshell Hume's conclusion means that we have no good reason to believe any of the information science gives us.

- Hume's conclusion is devastating for our everyday beliefs. Hume himself emphasizes that the results of his conclusion mean that we don't have good reason to believe that food and drink will nourish us rather than poisoning us—or even that the sun will rise tomorrow as it's always done in the past.

- Clearly, this is a terrible result. If there is any way to avoid accepting the force of Hume's argument, we should see if we can seek it out.

- One strategy for rejecting Hume's argument is to attack the first premise: that in order for you to have reason to believe any matter-of-fact claim about the unobserved, you must first have reason to believe the uniformity principle is true.

- This claim is not obvious. What is obvious is that for you to have good reason to believe the result of an inductive argument, the uniformity principle has to be true. But that's a far cry from Hume's claim, which is that for you to have reason to believe in the result of your inductive inference, you have to have *reason to believe* that the uniformity principle is true.

- If you're an externalist, you'll reject Hume's claim. Just because our having reason to believe in the conclusions of our inductive inferences requires the truth of the uniformity principle, that doesn't mean we have to give any thought to the uniformity principle itself.

So much for Hume's attack on induction. Unfortunately, we're not in the clear yet. In the 20th century, a number of philosophers returned to the question of why we're allowed to believe the conclusions of our inductive inferences.

READINGS

Stroud, *Hume*.

Swinburne, ed., *The Justification of Induction*.

QUIZ

1 Which of the following is not a way that induction differs from deduction?

- **a** Induction doesn't provide certain support for its conclusions.
- **b** The support that the premises of an inductive argument provide for the conclusion often involve probabilities.
- **c** Even if the conclusion of an inductive argument is well supported by the argument, the conclusion still allows for exceptions.
- **d** The conclusions of inductive arguments cannot contain more evidence than is contained in the premises.

2 TRUE OR FALSE

David Hume's uniformity principle, according to which unobserved instances resemble observed instances, lies at the foundation of the natural sciences.

3 David Hume's argument against the rationality of induction relies on which of the following?

- **a** The truth of internalism about knowledge
- **b** The idea that all claims involve either relations of ideas or matters of fact
- **c** The idea that induction relies on the uniformity principle
- **d** All of the above
- **e** None of the above

4 TRUE OR FALSE

A coherence theorist can accept all of the premises of David Hume's argument against the rationality of induction but still reject Hume's conclusion.

Answer key can be found on page 207.

LECTURE 15

The Raven Paradox and New Riddle of Induction

As introduced in the previous lecture, David Hume raised skeptical worries about the rationality of believing on the basis of inductive inference, and his argument—though ingenious—can be rejected in a variety of ways. There are problems involved with inductive inference that are not as easily solved as Hume's argument, but thinking about a solution to these problems provides valuable lessons about inductive inference.

HEMPEL'S PARADOX

- Hempel's paradox was formulated in the 1940s by German philosopher Carl Hempel. It is often referred to as the raven paradox because Hempel formulated it by discussing the claim that all ravens are black and considering what types of experiences would support that claim.

- Consider the claim that all ravens are black. If you want to find evidence for it, presumably the best way would be to find all the ravens you can and show that all the ravens you've been able to find so far have been black. Furthermore, it's intuitively plausible that the more black ravens you've found, the stronger your evidence is for the claim.

- Logically, the claim that all ravens are black is equivalent to the claim that all non-black things are non-ravens. "Equivalent" here means that the two statements are true in exactly the same circumstances.

- Whenever the "all ravens" claim is true, so is the "all non-black things" claim. And whenever the "all ravens" claim is false, so is the "all non-black things" claim. And this means that any evidence for the "all ravens" claim will also count, logically, as evidence for the "all non-black things" claim—and vice versa.

- The problem is that you can get a ton of evidence for the claim that all non-black things are non-ravens without ever seeing a single raven. Suppose, for example, that you're in a factory that mass-produces white printer paper. Each sheet of paper is both non-black and not a raven, so each sheet of paper provides evidence for the claim that all non-black things are non-ravens.

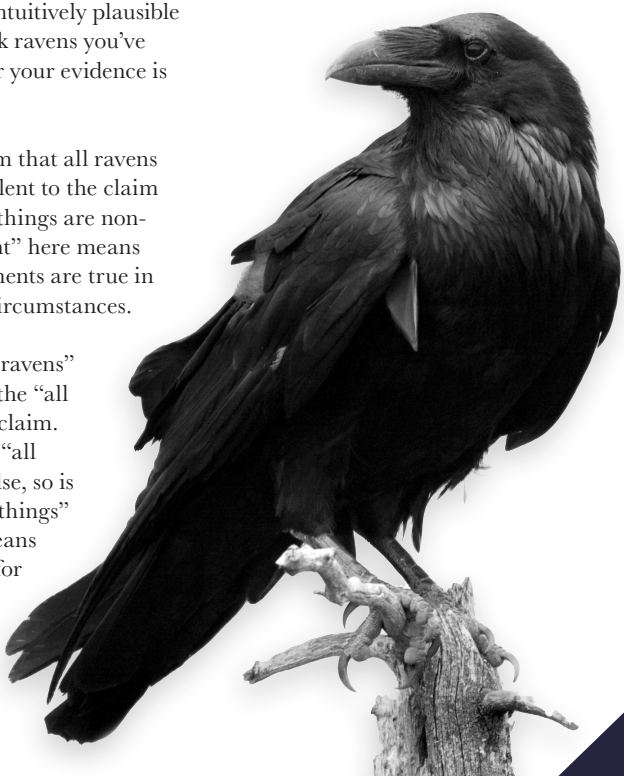

- But remember that the claim that all non-black things are non-ravens is equivalent to the claim that all ravens are black and that evidence for one claim should therefore count logically as evidence for the other. But this means that you can acquire a great deal of evidence in support of the claim that all ravens are black without ever laying eyes on a single raven!

- Also, remember that the more instances you have that support your claim, the stronger your evidence is supposed to be. So, it would seem that you can gather stronger evidence for the claim that all ravens are black by flipping through a ream of printer paper than you can by observing a few black ravens. This is absurd!

- The raven paradox demonstrates that there is more to the way that inductive arguments provide evidence for their conclusions than we can express using logic alone. Beyond this, though, philosophers disagree about the best way to deal with Hempel's raven paradox.

- One solution, proposed by philosophers Israel Scheffler and Nelson Goodman, drew on a famous insight of Karl Popper's. Popper argued that in science, evidence against a hypothesis, called disconfirmation, is much more important than evidence for that hypothesis, called confirmation.

- Scheffler and Goodman suggested that we shouldn't just focus on the fact that because the claim that all ravens are black and the claim that all non-black things are non-ravens are logically equivalent, both an instance of a black raven and an instance of a non-black non-raven would offer some confirmation for either claim. Instead, we should look at the claims that the instances of the black raven and the non-black non-raven can disconfirm.

- The instance of a black raven confirms the claim that all ravens are black, and it also provides reason to reject the claim that all ravens are non-black. The instance of a non-black non-raven, however, is compatible both with the claim that all ravens are black and with the claim that all ravens are not-black.

- Scheffler and Goodman characterize this difference by saying that the instance of a black raven selectively confirms the claim that all ravens are black, while the instance of a non-black non-raven doesn't selectively confirm that claim.

- Although we weren't able to deal with Hempel's paradox quite as easily as Hume's challenge to induction, there is nevertheless a pretty tidy solution that we can appeal to.

GOODMAN'S PROBLEM

- Goodman's problem for induction is often referred to as the new riddle of induction (the old riddle of induction is Hume's) or as the grue problem.

- Goodman begins by defining a description-word, "grue." The philosophical term for a description-word is predicate; it is basically a term used to describe something. For example, "tall" and "wears glasses" are predicates. And, after Goodman introduced the term in his 1955 work *Fact, Fiction, and Forecast*, so is "grue."

- Goodman defines grue by saying that

> *it applies to all things examined before [some future time]* t *just in case they are green but to other things [after that future time] just in case they are blue. Then at time* t *we have, for each evidence statement asserting that a given emerald is green, a parallel evidence statement asserting that that emerald is grue. And the statements that emerald* a *is grue, that emerald* b *is grue, and so on, will each confirm the general hypothesis that all emerald are grue. Thus according to our definition, the prediction is that all emeralds subsequently examined will be green and the prediction that all will be grue are alike confirmed by evidence statements describing the same observations. But if an emerald subsequently examined is grue, it is blue and hence not green.*

- Let's give a slightly more concrete example by defining grue as something that, whenever we examine it before midnight on January 1, 2100, appears green, but that were we to examine it on or after midnight on January 1, 2100, would appear blue.

- Because of the way we've defined grue, all of the experience that we've had so far of green things is also evidence that those things are grue.

- In particular, the following two statements are both true:
  - I All observed emeralds are green.
  - II All observed emeralds are grue.

- If observed correlations provide us with reasonable beliefs about unobserved cases, we have equal reason to believe the following two claims:
  - I′ The first emerald observed on or after midnight, January 1, 2100, will be green. (Let's call this the green-prediction.)
  - II′ The first emerald observed on or after midnight, January 1, 2100, will be grue. (Let's call this the grue-prediction.)

- But now we have a problem. Because of how we've defined grue, the grue-prediction implies the following:
  - II″ The first emerald observed after *t* will be blue. (Let's call this the blue-prediction.)

- Obviously, the blue-prediction contradicts the green-prediction.

- So, by introducing grue, Goodman demonstrates that it seems induction can be used to support outlandish beliefs—such as the belief that emeralds will appear blue when observed on or after January 1, 2100. And induction can even be used to support contradictory claims. This is because, by using grue, we've shown that both the claim "All emeralds will appear green when observed on or after January 1, 2100" and the claim "All emeralds will appear blue when observed on or after January 1, 2100" are equally well supported by the evidence.

- Goodman doesn't think we can find a solution to the grue problem by focusing on the logic of the predicates themselves. Instead, he thinks the

best way to distinguish between the predicates we want to employ and the ones we don't is to focus on our own psychology. How do we use predicates in formulating inductive arguments? Goodman thinks it's the answer to this question that provides the key to solving his new riddle of induction.

- Goodman observes that the claims we want to test through induction don't exist individually, in a vacuum. Instead, those claims are often related to a number of other claims. And one of the ways in which different claims are related to each other is that we often employ the same predicates to formulate them. Goodman describes those predicates that are used in a wide variety of claims that we test through inductive inference as entrenched.

- Now we have a way to distinguish between green and grue: Green is very well entrenched while grue is not entrenched at all.

- Green and grue conflict with each other because they support contradictory predictions. So, if we have to choose between two claims that we're testing against the evidence and one of those claims involves the predicate "green" while the other involves the predicate "grue," then Goodman says that we should choose the claim that involves the better-entrenched predicate, "green." In this case, the predicate "green" overrides the predicate "grue."

- This gives Goodman a way to characterize the kinds of predicates he thinks make good candidates for inductive inference—and offer a solution to the grue problem. It's the predicates that override other predicates that conflict with it without being overridden themselves. He calls those predicates projectible.

- Goodman's solution to the grue problem might seem too easy. What justification could we give for characterizing certain descriptions, such as "green," as projectible? Why are others, such as "grue," non-projectible? Wasn't the whole point of Goodman's puzzle that both descriptions fit our evidence equally well?

Bayes's Theorem

- One way to defend Goodman's type of solution is to see it as part of a larger theory of inductive inference. That's the strategy a number of contemporary philosophers pursue, by suggesting that the solution to Goodman's grue problem is just a special application of a rule of inference known as Bayes's theorem, named for the reverend Thomas Bayes, who formulated the theorem.

- The insight that Bayes captured was that when we learn new information about the probability of the occurrence of some event, we often already have information about how likely that event is. Bayes's theorem presents a precise mathematical formula for how to interpret the new information about likelihoods in light of the information that we already have.

- Suppose that your doctor suggested that you get tested for a very rare but very deadly disease. Let's say it's so rare that only one in 10,000 people ever get it. The test for the disease is 99 percent reliable. A few days later, you get a call from the doctor, saying that you have this rare deadly disease. How should you feel about your chances of actually having the disease?

- You might think that you have a 99 percent chance of having the disease. But you'd be wrong. And this is where Bayes's theorem comes in.

- Imagine that we have a random sample of 100,000 people. One in 10,000 of those people will have the disease, so that means that 10 will have it and 99,990 won't have it.

- The test is supposed to be 99 percent reliable. For simplicity's sake, let's say that means that for every 100 people who actually do have the disease, the test only gets it wrong one time—called a false negative—and that for

every 100 people who actually do not have the disease, the test again only gets it wrong one time—called a false positive.

- In other words, one out of 100 people that the test says doesn't have the disease actually does have it, and one out of 100 people that the test says does have the disease actually doesn't have it.

- That means that out of our sample of 100,000 people, the test will say that all 10 of those who do have the disease actually do have it. However, it will also—wrongly—say that one percent of those who don't have the disease actually do have it. That means it will falsely diagnose 1,000 people as having the disease when they actually don't have it.

- Let's put those numbers together. That 99-percent-reliable test diagnoses 1,010 people out of 100,000 as having the rare deadly disease, but only 10 of those people actually have the disease, so your chances of having the disease—even if that 99-percent-reliable test says you have it—are actually only 10 out of 1,010, or about one percent.

- This example shows the power of using what we already know about the world—called prior probability—to interpret any new information that we learn. In this example, the prior probability has to do with just how very infrequently that deadly disease occurs. Without taking the prior probability into account, when you get the bad news of the test results for the disease, you would think that you are almost certainly a goner. Actually, though, 99 out of 100 people who get similar bad news go on to live a long life.

- A number of philosophers have suggested that Bayes's theorem is the basis for the correct understanding of how to deal with Goodman's grue problem. In fact, many philosophers think that by applying the insights that Bayes introduced, we may solve a number of problems in both epistemology and the philosophy of science.

We've seen reasons for thinking that induction can in fact be a powerful tool for acquiring new information. This tells us that internalist philosophers who develop rules for good inferential practices are making important contributions to our understanding of induction and how to improve it.

However, we've also seen reasons for thinking that, as with deduction, human beings are just not very reliable inductive reasoners. This suggests that if we're trying to evaluate the structure of human knowledge, externalism still seems like the best theory for explaining the contributions of inference to knowledge.

READINGS

Skyrms, *Choice and Chance*.

Stalker, ed., *Grue*.

QUIZ

1 Which of the following is equivalent to the claim that all NBA players are tall?

- **a** Everyone who is tall is an NBA player.
- **b** All non-NBA players are non-tall.
- **c** All non-tall things are non-NBA players.
- **d** All of the above.
- **e** None of the above.

2 TRUE OR FALSE

According to Karl Popper, confirmation is more important than disconfirmation.

3 TRUE OR FALSE

According to Israel Scheffler and Nelson Goodman, the solution to Carl Hempel's paradox relies on the fact that the existence of a black raven not only provides evidence for the claim that all ravens are black but is also incompatible with the claim that all ravens are non-black.

4 TRUE OR FALSE

A solution to Nelson Goodman's new riddle of induction is that "grue" is a time-indexed predicate whereas green is not.

5 According to Bayesian reasoning, if a person who is almost always reliable—say 99 percent reliable—tells you that a one-in-a-million event occurred, then the chance that the event actually occurred is which of the following?

- **a** 99 percent
- **b** Greater than 99 percent
- **c** Less than 99 percent

Answer key can be found on page 207.

LECTURE 16

Know-How versus Propositional Knowledge

In addition to propositional knowledge, another type of knowledge is know-how. The objects of know-how are of a different type than those of propositional knowledge. Rather than taking as its object a proposition, or the sort of thing that could serve as the content of a claim, know-how involves a skill or performance.

PERFORMATIVE VERSUS ACQUAINTANCE KNOW-HOW

- Consider the following cases:
  - Steve is a tennis coach who knows how to serve a tennis ball over 120 miles per hour, though he is unable to do so that fast himself.
  - Carl knows how to tell the sex of baby chicks, though he does not know how he knows.

- Each of these is a case in which we naturally say that Steve or Carl possesses know-how. They each differ in a variety of ways, but there are also a number of similarities—two of which are particularly significant.

- First, both cases involve a reliable disposition or ability on the part of Steve and Carl. Here, a disposition is simply a tendency to act, speak, or believe as long as certain conditions are met; this is important because a person may have certain dispositions that never become apparent.

- This explains the case of Steve. Although he cannot serve the ball over 120 miles per hour himself, his knowledge of how to serve the ball that fast manifests itself, for example, in his disposition to tell his tennis students how to alter their serve mechanics so that they can improve their service speed.

- Carl doesn't have a mere disposition; he is actually able reliably to distinguish male from female chicks.

- In both cases, we attribute the disposition or ability to the people performing the actions—to Steve or Carl.

- Recall that when discussing self-knowledge, we distinguished between computations or inferences that *you* perform and computations or inferences that *your brain* performs, without any awareness or intention on your part. That distinction applies again here. Neither of these cases involves mere actions of Steve's or Carl's bodily organs or muscles.

- In addition to these similarities among the example cases, there are some important differences. In particular, while Carl's knowledge of how to sort baby chicks involves a reliable ability to sort those chicks, Steve's knowledge of how to serve a tennis ball more than 120 miles per hour does not involve a reliable ability to serve a ball that fast. While Steve can recognize the mechanics needed to achieve that velocity on serve and can tell his students how to achieve those mechanics, he isn't able to replicate the mechanics himself.

- To distinguish Carl's know-how from Steve's, we can say that Carl has performative know-how.

- There is a second major difference between Steve's know-how and Carl's. It seems plausible to explain know-how like Steve's in terms of a reliable recognitional ability. When he perceives different mechanics of serving, he is reliably able to say which will result in a serve of at least 120 miles per hour. When we discussed different kinds of knowledge, we characterized this type of ability to recognize something in perception as acquaintance knowledge. For this reason, we can say that Steve has acquaintance know-how.

- In some cases, people can have both performative and acquaintance know-how. In fact, in some cases, it would probably be correct to explain someone's performative know-how in terms of their acquaintance know-how. But it's also possible for somebody to have performative know-how without having acquaintance know-how. In fact, Carl is an example of that. Although some part of Carl's brain performs calculations based on a certain way the chicks appear, Carl himself is not acquainted with any particular way of distinguishing between male and female chicks. He just knows that he's good at it.

KNOWLEDGE-HOW VERSUS KNOWLEDGE-THAT

- There is some parallel between the performative sense of knowledge-how and the externalist, reliabilist view of knowledge-that. Carl, for example, is reliable at distinguishing male from female baby chicks, without knowing how he's able to be so reliable.

- Analogously, there are some parallels between the acquaintance sense of know-how and internalist views of knowledge-that. Steve's know-how is closely tied to his having a representation of the kind of mechanics necessary to serve a tennis ball over 120 miles per hour. And Steve is able to verbalize his knowledge in a way that is helpful to the players he coaches.

- Relatedly, internalism about knowledge-that supports the idea that the justifications of our knowledge are the sorts of things we're often aware of. Because we're aware of our justifications, we can not only verbalize them and transmit our knowledge to others but also explain to them the reasons for our knowledge.

- British philosopher Gilbert Ryle was the first to argue explicitly for the claim that know-how is a form

of knowledge utterly distinct from factual or propositional knowledge—knowledge-that. He suggests that there's a threat of infinite regress if we think all knowledge is knowledge-that, or propositional knowledge. Because that would be bad, Ryle thinks we're forced to acknowledge that there's another type of knowledge that is different in kind from knowledge-that. He thinks this different type of knowledge is knowledge-how.

The expression "the ghost in the machine" is one of Ryle's colorful ways of dismissing the idea of internal, unobservable mental states.

In his book *The Concept of Mind*, Ryle argues that knowledge-how is a form of knowledge utterly distinct from knowledge-that:

> *To put it quite generally, the absurd assumption made by the intellectualist legend [i.e., the idea that all knowledge-how is a form of knowledge-that] is this, that a performance of any sort inherits all its title to intelligence from some anterior internal operation of planning what to do. … By the original argument, therefore, our intellectual planning process must inherit its title to shrewdness from yet another interior process of planning to plan, and this process could in its turn be [analyzed with respect to its shrewdness]. The regress is infinite.*

- In our discussion of Steve and Carl, we saw a number of independent reasons for thinking that at least one sense of knowledge-how, the performative sense, is not a form of knowledge-that. This would seem to provide some support for Ryle's claim.

- An influential challenge to Ryle was published in the *Journal of Philosophy* in 2001 by philosophers Jason Stanley and Timothy Williamson, who read Ryle as trying to show that the assumption "all knowledge-how is really just propositional knowledge" leads to a contradiction. Ryle does this by suggesting, first, that if you perform some action intentionally, then you're employing knowledge-how to perform that action. From there, he suggests that if you employ some of your knowledge-that, or propositional knowledge, it means that you must be contemplating the proposition that you know.

- Here, according to Stanley and Williamson, is how Ryle thinks the problematic regress gets started. According to this reading of Ryle, everyone who performs some action intentionally is employing knowledge-how. And if the "intellectualist legend" is true, this means that by employing knowledge-how, they're actually employing knowledge-that.

- Everyone who employs knowledge-that, according to this reading of Ryle, must contemplate the proposition that they know. Contemplating a proposition, however, is itself an action. And now we're trapped in the beginning of the regress.

- Stanley and Williamson object to Ryle's argument. To do so, they appeal to an earlier reply to Ryle given by philosopher Carl Ginet in his book *Knowledge, Perception, and Memory*. They argue that Ryle is making a mistake if he's suggesting that whenever anyone employs propositional knowledge they're also performing the distinct action of contemplating the proposition that they know. Why think that contemplating a proposition is itself a distinct and separate action?

- That does indeed seem plausible. And the only way Ryle can dispute Ginet's argument, Stanley and Williamson suggest, is by interpreting the contemplation of propositions as a nonintentional—or otherwise deflated sense—of action. But the whole point of Stanley and Williamson's reading of Ryle's argument requires that the actions we're looking at be intentional actions. Without that, the regress won't get off the ground.

- In other words, according to Stanley and Williamson, Ryle needs two fundamental claims to make his infinite regress argument against the intellectualist legend: Ryle needs it to be the case that when you contemplate a proposition, you're performing an

action, and he needs it to be the case that it should count as an intentional action. But those two things together, according to Stanley and Williamson, just aren't plausible. So, it seems that Ryle's argument just won't work.

- If Stanley and Williamson are reading Ryle correctly, then it seems that Ryle's argument rests on a mistake. The regress that Ryle tries to use against what he calls the "intellectualist legend" can't even get started.

UNDERSTANDING RYLE'S ARGUMENT

- Stanley and Williamson take Ryle to be talking about intentional actions. But in the passage, Ryle never uses the words "intend," "intention," or "intentional." Instead, he talks about "intelligent" actions.

- Stanley and Williamson do Ryle a disservice by taking his argument to be dealing with all intentional actions, more broadly. Rather, Ryle is focusing on the narrower category that he refers to as "intelligent."

- In the context of his discussion, Ryle is using the term "intelligent" to characterize the actions that display knowledge-how. This also makes more sense. Stanley and Williamson's version of Ryle's argument relied on the claim that every intentional action involves knowledge-how. But that's just not a plausible claim.

- Here's why it's not plausible to think that every action that you intend involves know-how. When you're learning how to perform some action skillfully, you intend to perform the action, but you don't yet know how to do it. That's the only way you can acquire know-how in the first place.

- Here's how we ought to understand Ryle's argument.

- First, we should characterize the "intellectualist legend" that Ryle is attempting to criticize. According to the "intellectualist legend," all knowledge-how is a type of knowledge-that. The intellectualist presumes that whenever you perform an action on the basis of your knowledge-how, there is a certain way you have for performing that action—let's call that your method. With that assumption, we can characterize the "legend" as the thesis of intellectualism, like this:

> *Intellectualism says that know-how involves guiding your performance on the basis of your knowledge of a proposition. The proposition you know is that your method is a way to perform the action.*

- The argument then seems to run as follows:
  1. If your action is a knowledgeable one, then you perform that action on the basis of your knowledge-how.
  2. Intellectualism is true.
  3. So, then you must have a method on which your performance of the action is based, and you must guide your performance on the basis of your knowledge of the fact that your method is a way to perform the action.
  4. If you guide your performance on the basis of your propositional knowledge, then that guidance itself must be knowledgeable.

- But now we have to apply the intellectualism thesis again, to the fourth premise, and so on. That looks like an infinite regress.

- This version of Ryle's argument does not rest on any independently implausible theses concerning what is involved in employing knowledge-that, or propositional knowledge. In fact, it seems to employ only claims that a defender of the intellectualist legend would accept.

- The intellectualist would not object to the first claim. All it does is capture the difference between action on the basis of knowledge-how—i.e., knowledgeable action—and action that accidentally achieves the desired result.

- It also doesn't seem that the intellectualist can reject the thesis of intellectualism, because the intellectualist will need that thesis to rule out the types of cases, like Carl, that motivated us to introduce the performative sense of knowledge-how in the beginning of the lecture.

- This is because for the intellectualist view of knowledge-how to work, intellectualists would have to claim that the only sort of knowledge-how is captured by the acquaintance sense of knowledge-how. Then, they could say this: When you perform an action on the basis of know-how, it's because you're acquainted with a method for performing that action and you know that the method you're acquainted with is a way to reliably perform the action. That's just what it means to say that knowledge-how is a species of knowledge-that, or propositional knowledge. And in fact, that's exactly Stanley and Williamson's view.

- That's why it looks like the opponents of Ryle's position would have to accept the intellectualism thesis—and the second step of the argument.

- The next step of the argument we have to consider is the fourth claim. Why think that your guidance of your performance also has to be knowledgeable? The reason is that there has to be a connection between the acquaintance knowledge you have and your performance of the action.

Otherwise, you wouldn't actually be performing the action on the basis of your knowledge-how.

We've covered all of the steps in the argument that the intellectualist could possibly try to criticize. It seems that Ryle's infinite regress argument is a strong one after all. And this means that there is a type of know-how, such as the procedural know-how possessed by Carl, that isn't a form of propositional knowledge.

As with our discussions of externalism, the discussion of the performative sense of know-how suggests that it would be a mistake to place too much emphasis on the role of our internal mental states in knowledge.

READINGS

Bengson and Moffett, eds., *Knowing How*.

Stanley, *Knowledge and Practical Interests*.

QUIZ

1 H. M. had no long-term memory but was taught the mirror-drawing skill. H. M. doesn't know that he has that skill and doesn't remember learning the skill. Plausibly, we can describe his know-how as which of the following?

a Acquaintance know-how

b Performative know-how

2 Suppose that you are a golf coach who knows a great deal about golfing mechanics, but because of nerve damage, you are no longer able to hold a golf club. Your ability to instruct your students on putting technique can plausibly be described as which of the following?

a Acquaintance know-how

b Performative know-how

3 TRUE OR FALSE

According to Jason Stanley and Timothy Williamson, the root of the problem with Gilbert Ryle's argument against the equivalence of knowledge-how and knowledge-that is that Ryle mischaracterizes the nature of intentional action.

4 According to a more charitable reading of Ryle's argument, which of the following is not true of that argument?

a It suggests that we must distinguish knowledge-how from knowledge-that to avoid a vicious infinite regress.

b It suggests that intelligent action involves knowledge.

c It suggests that all actions are intentional actions.

Answer key can be found on page 207.

LECTURE 17

Knowledge Derived from Testimony

Most of the discussions in this course so far have centered around individualistic sources of knowledge; in other words, the sources have involved skills or abilities that you could exercise in solitude. However, much of what you know was gained by means of consulting others' expertise. In philosophy, the process of acquiring information through language-based communication with others is called testimony. Consider all of the ways in which testimony is woven into the fabric of our knowledge, from large-scale science and team-based medical research to investigative reporting—any of which would be unthinkable without relying on information derived from testimony.

NON-PRESUMPTIVISM

- Suppose we grant that testimony is an extremely important source of knowledge. We're then faced with the question of how the beliefs we form on the basis of testimony achieve the level of knowledge. Someone tells you something, or you read something, and you believe it. Is that enough to count as knowledge if what you believe turns out to be true? How does that work?

- One explanation was given by David Hume:

> *[A]s the evidence, derived from witnesses and human testimony, is founded on past experience, so it varies with the experience, and is regarded either as a proof or a probability, according as the conjunction between any particular kind of report and any kind of object has been found to be constant or variable.*

- In other words, it is not enough simply to be told something or to read something in order to have good evidence. Instead, if you believe something on the basis of someone's testimony, you have to support your belief on the basis of an argument.

- This may seem surprising, given that the last time we encountered Hume he was arguing that inductive reasoning never provides us with actual reasons for our beliefs. However, when it comes to testimony, Hume suggests that the evidence we derive from testimony is simply a form of inductive inference.

- Suppose you receive testimony from Tim. Hume thinks that you ought to reason according to something like the following basic structure:
  1. Tim, who has certain attributes that are relevant to you taking him seriously in this kind of situation, gives you some piece of information.
  2. Generally, when speakers who are similar to Tim—and in similar types of situations—give you some piece of information, their reports are reliable.
  3. Thus, (probably) the information Tim has given you is likely true.

Say you're a detective and you're trying to gather information about the location of a suspect at the time a crime occurred. You ask the suspect's secretary, who doesn't particularly like the suspect and who has no knowledge of when the crime occurred, and she gives you information about the suspect's activities, from which you gather that the suspect was nowhere near the location of the crime at the time the crime occurred.

In this case, the relevant attributes of the testifier include her lack of motivation to lie to give an alibi to the suspect as well as her lack of knowledge about the time for which the suspect would need an alibi. In a situation like this one, you can reasonably assume that the secretary is telling the truth and conclude that the suspect probably has a strong alibi for the time in question.

When you consider examples like this, Hume's view about testimony sounds pretty strong.

- Because Hume requires that you have positive reasons, in the form of an argument, for accepting that testimony as evidence for knowledge, we can call his view inferentialist non-presumptivism. It's inferentialist because you need to support your acceptance of testimony by means of an inductive inference. And it's non-presumptivist because testimony doesn't enjoy any presumption of truth; you have to have positive reasons for accepting someone's testimony if you're going to rely on that testimony as evidence for knowledge.

- Even though Hume's view is one of the most popular positions historically about the nature of evidence that testimony provides for knowledge, it's not the only influential theory.

- Pretty soon after Hume proposed the theory, other philosophers noted problems with it. One of the strongest critics of Hume's inferentialism was philosopher Thomas Reid, who raised a number of criticisms of Hume's theory of testimony—of which we will focus on two.
  - Reid criticizes Hume for getting the phenomenology of believing on the basis of testimony wrong. Normally, when someone tells you something and you don't have a particular reason not to believe the person, then you just do believe them, without rehearsing any sort of argument in support of your belief.
  - Reid asks us to consider how children learn from the people around them and suggests that children aren't capable of using inductive inference to support their reliance on other people's testimony. Not only do they lack the reasoning skills, but they also lack the experience needed to have a basis for an inductive argument. If Hume were correct, Reid says, you would expect children to be the most mistrustful of testimony, while the most knowledgeable, most experienced people would be the most trusting of testimony. As Reid notes, however, this is not the case. Children are quite willing to accept the testimony of the people around them.

- Reid defends a rival theory about the way that testimony provides support for knowledge. Unfortunately, it is not entirely clear what theory Reid advances to replace Hume's. We'll consider two options.

- Both rival theories to Hume share one big similarity: They assert that, at least sometimes, you can know something on the basis of someone's testimony even if you aren't aware of any positive reasons in support of your belief of that person's testimony. In other words, they reject Hume's requirement that you can't know something on the basis of testimony unless you rehearse positive reasons for accepting that testimony.

- Where the two challengers to Hume differ, however, is in their explanation for why they reject Hume's positive reasons requirement.

EXTERNALISM

- One motivation for rejecting Hume's positive reasons requirement comes from externalism. An externalist theory of knowledge on the basis of testimony would allow that if you are in circumstances in which the testimony you're receiving is a reliable source of information, you can

acquire knowledge on the basis of that testimony as long as the reason you believe it is that your acceptance of the testimony involves a reliable method.

- Most externalist theories of testimonial knowledge think that the phrase "your acceptance of the testimony involves a reliable method" involves at least two components.
  - You can't have any positive reason to question the circumstances or the speakers. If you're in a used car showroom and you know that the salesperson earns most of his or her money from commissions, then you shouldn't automatically trust the salesperson when he or she says the car you're interested in has never been in any major accidents.
  - You also have some set of unconscious abilities that monitor the trustworthiness or honesty of your sources of information. In face-to-face communication, this would include signals the speaker gives off that you might not even be consciously aware of but that are supposed to alert you if your brain is triggered by anything suspicious. Let's call this the unconscious monitoring requirement.

In our discussions of perception and memory, we've seen that externalists allow for the possibility that we might not know when our perceptual abilities are reliable but that we can still know information we acquire using those abilities.

For the sort of externalist theories we've been considering, it is enough that the methods we use to acquire information are in fact reliable. Those theories don't require that we also have good evidence that they are reliable.

Our culture is full of evidence that many of us at least intuitively accept the idea of an unconscious monitoring requirement. For example, you might read an interview with a policeman who caught a serial killer during a routine traffic stop and perhaps the policeman says, "Something just seemed off about the guy when I stopped him."

- An advantage of the externalist take on testimony is that it can treat testimony as completely on par with perception or memory. As with those sources of knowledge, we are not always aware when we are reliable or what the mechanism is that ensures we are reliable. What is important in all cases, for the externalist, is the fact *that* we are reliable.

- Hume's non-presumptivist inferentialist theory differs from externalism, then, in one very important respect: It requires you to have a positive argument in support of your acceptance of someone's testimony. In contrast, the externalist theory only requires that someone's testimony actually is a reliable source of information if you are going to rely on that testimony.

- However, like Hume, the externalist also thinks you have to be sensitive to relevant features of the testimony that serve as indicators of its reliability. It's just that the externalist thinks that you can be sensitive to these features unconsciously and that you don't have to use them to formulate an explicit argument in support of your acceptance of a given piece of testimony.

PRESUMPTIVISM

- Another rival view agrees with the externalist in rejecting Hume's positive reasons requirement but goes even further by also rejecting the externalist requirement of unconscious monitoring of testimony for indicators of reliability. Let's call this view presumptivism, because the various versions of the view all involve a presumptive right to accept someone's testimony.

- These views hold that whenever you receive a piece of testimony, you have the right to assume, without rehearsing positive evidence, that what the testimony says is true, unless you have specific evidence that overrides that presumption.

- Presumptivism doesn't seem to have anything to do with knowledge. In order for a source of information to be a source of knowledge, it should reliably provide you with accurate information. Hume's theory of testimony accounts for this by requiring you to give an explicit argument that the testimony you're receiving is reliably tied to the truth. And externalism accounts for this by requiring that the testimony in fact be reliably accurate. But presumptivism doesn't seem to do anything to account for the fact that we need sources of knowledge to be reliable sources of accurate information.

- The standard strategy for defending presumptivism is to argue that anything more than presumptivism would actually be requiring too much and would make knowledge from testimony too hard to come by.

- One version of this defense of presumptivism argues simply that presumptivism is the only explanation available for how we gain knowledge through testimony. Another version argues in particular that without accepting presumptivism, we would have no explanation for how we communicate with others using language. Because language is required to gain knowledge from testimony, however, this version is also ultimately arguing that the only explanation for how we gain knowledge from testimony requires that we accept presumptivism.

- While there doesn't seem to be a good reason to accept presumptivism, there are good reasons to reject it. In particular, it doesn't do enough to ensure that the testimony we rely on will actually be accurate enough to support knowledge.

- Although Hume's inferentialism does better than presumptivism in this regard, there are some good reasons not to accept it. In particular, it doesn't do a good job of accounting for how we seem to deal with other people's testimony or of explaining how little children learn from testimony.

- That leaves us with externalism. Like presumptivism, it doesn't require that we give an argument in support of our acceptance of testimony, but like Hume's non-presumptivism, it does require that we have some way to ensure that the testimony we rely on will be reliably accurate.

READINGS

Coady, *Testimony*.

Lackey and Sosa, eds., *The Epistemology of Testimony*.

QUIZ

1 David Hume's insistence that you can only have knowledge through testimony if you support your reliance on that testimony on the basis of an argument that the testimony is accurate is an example of which of the following?

a Inferentialist non-presumptivism
b Presumptivism
c Externalism

2 If Thomas Reid thought that you can have knowledge through testimony so long as the testimony is accurate and you didn't have any reason to doubt that testimony, then that would be an example of which of the following?

a Inferentialist non-presumptivism
b Presumptivism
c Externalism

3 If Thomas Reid thought that you can have knowledge through testimony only if you rely on—perhaps unconscious—faculties that assess that testimony for accuracy, then that would be an example of which of the following?

a Inferentialist non-presumptivism
b Presumptivism
c Externalism

4 Which of the following is not a reason for doubting inferentialist non-presumptivism?

a Very young children can acquire knowledge through testimony.
b It often seems as if we accept testimony without rehearsing an argument in support of that acceptance.
c David Hume didn't think that inductive arguments provide us with reasons for belief.

Answer key can be found on page 207.

LECTURE 18

Social Psychology and Source Monitoring

In the previous lecture, we discovered that neither presumptivism nor Hume's non-presumptivism was very promising as a theory of how we acquire knowledge through our belief in testimony. It looked like externalism's notion of unconscious monitoring offered the best of both worlds: It allowed us to avoid having to follow Hume and require an explicit argument in support of testimony, something that didn't seem very plausible, and it also allowed us to avoid having to follow the presumptivist and advocate what looked like excessive gullibility, something that didn't seem very plausible either. However, the unconscious monitoring relied on by the externalist theory won't work either. There is a great deal of evidence that we try to look out for signs that the testimony we rely on is accurate—in fact, few results in social psychology are better supported—but we're just not very good at monitoring for signs that the testimony we rely on is accurate.

MONITORING FOR SINCERITY AND DECEPTIVENESS

- More than 50 years of research has established that humans are quite unreliable at responding to indicators of trustworthiness or deceit. In fact, most people would do just as well to guess or flip a coin to determine whether someone is lying or being truthful.

- One reason why we're so unreliable at detecting truthfulness has to do with the signs that we look for when we're judging whether someone is sincere. Some of those signs include maintaining good eye contact, not smiling or laughing too much, and projecting confidence—all of which we rely on. The problem is that none of those signs are actually reliable indicators of sincerity.
  - Repeated research since the early 1980s has demonstrated that people who are attempting to deceive you actually sustain eye contact better than people who are being sincere.
  - Research has also shown that when it comes to keeping a straight face, people who are being deceptive don't have more nervous smiles than truth tellers. If anything, liars smile somewhat less often.
  - In one of the most surprising studies of confidence, patients were asked to interact with doctors—actually, actors playing doctors. One actor was told to act blithely confident, despite admitting that the area wasn't an area of his expertise. The other actor was told to admit that the area wasn't one in which he was an expert and, in front of the patient, consult a reputed textbook on the subject. Patients preferred the blithely confident "doctor" to the one who admitted his lack of expertise but who cited the appropriate source in the presence of the patient.

- So much for trustworthiness. Perhaps, however, our unconscious monitoring is directed more at warning us when someone is being deceptive rather than at reassuring us when someone is being sincere. Is that a promising route to defending our unconscious monitoring abilities?

- Unfortunately, decades of empirical research provide no empirical support for the notion that subjects are reliable at detecting deception. Here is how two of the most prominent researchers in the field, Paul Ekman and Maureen O'Sullivan, put it:

> *In every study reported, people have not been very accurate in judging when someone is lying. ... Average accuracy in detecting deceit has rarely been above 60% (with chance being 50%), and some groups have done worse than chance.*

- Even people whose careers—if not their lives—depend on their ability to detect deceptiveness are not overall particularly reliable at it. Research conducted on customs officials, federal law enforcement officers, and police officers suggested that their overall reliability was so disappointing that, as one team of researchers concluded, "it is unlikely that judging deception from demeanor will ever be sufficiently accurate to be admissible in the courtroom" and "most of us would do well to entertain some skepticism about our ability to detect deception from demeanor."

- Maybe the problem is that the research so far has had to do with people trying to detect deceptive behavior in strangers. Maybe at least we get better at detecting deception in people we know. Here's the problem with that line of thought: Just as you get more familiar with someone and know more about the person's demeanor and so on, he or she is also getting more familiar with you and more able to tailor his or her way of communicating to your particular personality.

- Two of the most prominent researchers on deception detection, David Buller and Judee Burgoon, have suggested that "deceivers in interactive contexts should display increasing immediacy and involvement, pleasantness, composure, fluency, and smooth turn taking over the course of the interaction."

- And in fact, this is what studies seem to suggest. For example, a study conducted by Steven McCornack and Malcolm Parks in the 1980s examined partners in intimate relationships. Though the partners actually had greater confidence in their ability to detect deception, they were in fact worse at detecting deception in their partners than in others.

MONITORING FOR COMPETENCE

- Although we can't monitor for sincerity or deceptiveness, what about monitoring for competence? Surely we're better at recognizing when the people we're speaking with actually know what they're talking about, right?

- Unfortunately, the news about monitoring for competence isn't very good either. Consider a study conducted in the 1970s by Donald Naftulin and his colleagues, who gave the name Dr. Myron L. Fox to a television actor to pose as an authority on the application of mathematics to human behavior. The actor was then coached to present an hour-long talk on "Mathematical Game Theory as Applied to Physician

Education," full of nonsense, made-up terms, bad reasoning, and contradictions. Naftulin and his associates then examined the responses to that talk from three separate groups.

- The first group was composed of psychiatrists, psychologists, and social-worker educators attending a conference for educators of health professionals. They attended the talk and then had a half-hour discussion period during which they had an opportunity to question "Dr. Fox" further.
- The second group consisted of mental health educators—psychiatrists, psychologists, and psychiatric social workers. The second group, however, merely viewed the lecture and the discussion period on videotape.
- The third group was made up of educators and administrators taking a graduate-level education course. That group also viewed the lecture and discussion on videotape.

- In all of the groups, a vast majority of participants responded positively to the following questions about "Dr. Fox" and his lecture:
  - Did he use enough examples to clarify his material?
  - Did he present his material in a well-organized form?
  - Did he stimulate your thinking?
  - Did he put his material across in an interesting way?
- Some of the terms used in their feedback included "articulate," "knowledgeable," "good analysis of subject that has been personally studied before," "lively examples," "excellent presentation," "enjoyed listening," and "too intellectual a presentation."
- Naftulin and his colleagues reference Neil Postman and Charles Weingartner's claim in their book *Teaching as a Subversive Activity* that

> *it is the sign of a competent crap detector that he is not completely captivated by the arbitrary abstractions of the community in which he happened to grow up [when they conclude that] the three groups of learners in this study, all of whom had grown up in the academic community and were experienced educators, obviously failed as "competent crap detectors" and were seduced by the style of Dr. Fox's presentation. Considering the educational sophistication of the subjects, it is striking that none of them detected the lecture for what it was.*

- The problem is that it is taxing to think too carefully about whether someone is competent. It is easier to rely on superficial signs of

competence—such as whether Myron L. Fox is called "doctor," is dressed professionally, or seems competent. Noted social psychologist Robert Cialdini calls this the "click, whirr mode" of reacting.

- Examples of the "click, whirr mode" abound. For example, in a 1996 study, Joel Cooper and his colleagues had a mock trial jury listen to one of two expert witnesses. The first expert witness, who was introduced as having outstanding credentials, gave a jargon-filled, incomprehensible testimony. The other expert witness gave exactly the same testimony but was introduced as having shaky credentials. The mock trial jury was much more likely to believe the testimony given by the expert witness introduced as having the outstanding credentials—even though the testimony itself was exactly the same jargon-filled nonsense.

- Surprisingly, when the witness was introduced as having outstanding credentials, the jury found him almost twice as persuasive when he presented his argument incomprehensibly than when he presented the same testimony using terms that the layperson could follow. In other words, the jury found the expert much more convincing when they had no idea what he was talking about!

- The way Cooper and his colleagues made sense of this was that the members of the mock trial jury actually preferred using the easier "click, whirr mode" of just deferring to the supposed expert's competence than evaluating his arguments on their merits.

None of this should make us very optimistic that we're generally all that reliable at unconsciously monitoring competence, either. So, in other words, we're three for three: bad at monitoring for sincerity, bad at monitoring for deceptiveness, and bad at monitoring for competence.

READINGS

Gelfert, *A Critical Introduction to Testimony*.

Shieber, *Testimony*.

QUIZ

1 TRUE OR FALSE
Evidence from psychologist G. Stanley Hall suggests that Laura Bridgman was not capable of the sort of monitoring of testifiers that externalism requires.

2 The case of Laura Bridgman provides the most problems for which of the following?

a Inferentialist non-presumptivism

b Presumptivism

c Externalism

3 TRUE OR FALSE
If you pay attention to signals like eye contact and demeanor, you can tell when someone is lying.

4 Because of their reliance on what Robert Cialdini calls the "click, whirr mode," jurors have shown that they prefer expert testimony from which of the following?

a Arguments that the jurors themselves can follow given by experts with poor credentials

b Jargon-filled, incomprehensible arguments given by experts with poor credentials

c Arguments that the jurors themselves can follow given by experts with impressive credentials

d Jargon-filled, incomprehensible arguments given by experts with impressive credentials

Answer key can be found on page 207.

LECTURE 19

Testimony through Social Networks

The previous two lectures have examined the three most prominent theories of how we can support our knowledge on the basis of the evidence we get from the testimony of other people: presumptivism, non-presumptivism, and externalism. But none of these theories seems adequate. There is an alternative picture that offers a new solution. Like externalism and presumptivism, this new solution rejects Hume's requirement that we support our belief in testimony on the basis of an explicit argument. Like externalism and Hume's non-presumptivism, this solution also suggests there is an important insight in requiring that there be a mechanism to ensure that the testimony we're relying on is reliably accurate. Unlike externalism, however, this new solution doesn't require that we base that assurance on any of the unreliable mechanisms social psychologists suggest we employ when we try to monitor testimony for signs of accuracy.

DUNBAR'S NUMBER

- One of the most famous numbers in social science is Dunbar's number, named after anthropologist and evolutionary psychologist Robin Dunbar, who first proposed it in the 1990s. As a researcher working with nonhuman primates, he was curious about why the primates he was studying spent so much of their time and effort grooming each other.

- In the 1980s, the view that came to be known as the social brain hypothesis was starting to gain attention in primate studies. According to this hypothesis, we can explain the comparatively large size of primate brains because of the social complexity of the groups they live in.

- Dunbar's idea was to link the size of a primate's social group—which he defined in terms of its grooming behavior—and the size of the part of the primate brain that is responsible for tracking complex information, the neocortex.

- From his grooming data, Dunbar figured out the average social group sizes for the primate species he was studying. From there, he hypothesized that the number of group members an average primate of that species would have in its social group could be predicted based on the ratio of the volume of the neocortex to the total brain volume of that species.

- But what sparked widespread interest in Dunbar's work was that he extrapolated his results to humans. Building on the ratio that he had established and based on the size of an average human brain, Dunbar predicted that the number of people that the average human has in his or her casual social group is around 150—the most famous of a series of Dunbar's numbers.

- Dunbar suggested that we each have different-sized social groups relating to each other by a precise formula that we can approximate with a rule of three. As we consider the people we have more distant connections with, we each have about 500 acquaintances and up to about 1,500 people that we can recognize by face.

- At the other extreme, we're most closely connected to about 5 people—our closest friends or family—then 15 people—the ones that we would confide in or turn to in a time of need—and then about 50 people who are our friends but not ones we would think to call in an emergency.

- The midpoint of those more intimate and more distant social relations is the Dunbar number that most people know: 150. Dunbar found that you didn't have the same 5, or 15, or 50 people in those different social groups but that the people who make up the groups can change over time. However, the number of people in the different groups stays roughly the same.

Years of research have confirmed the validity of Dunbar's numbers.

For example, armies throughout history have been organized roughly along the lines predicted by Dunbar's research. The smallest units tend to be between 10 and 15 people, with larger organizational units ranging up to companies of around 150 soldiers and up to battalions or regiments of between 500 and 1,000 personnel.

- Two of the scholars who have done the most to popularize the study of how the structure of our social networks influence different aspects of our lives are Nicholas Christakis and James Fowler. In their book *Connected*, they trace a number of effects that social networks can have on us, from our health to our political affiliations to our success in finding work. One of those effects is how we can gain information through our social connections.

WEAK TIES

- Sociology professor Mark Granovetter has emphasized the importance of so-called weak ties, which are our connections to the friends of our friends, or even to our friends' friends' friends.
- A simplistic picture of influence would be to think of the spread of information between people like a bucket brigade. You pass information to your friend, who passes information to his or her friend, and so on. The problem with this is that if we only pass information to those we're closely connected with—our strong ties—the information would never travel very far.
- And this is where the power of weak ties comes in. In an example from Christakis and Fowler's *Connected*, almost 40 percent of a piano teacher's recommendations weren't from direct acquaintances but from acquaintances three degrees removed.

NETWORK EFFECTS

- Not all social networks select the best ideas. Instead, the makeup of a social network and which components of that network have influence can affect which ideas come out on top in that network.
- This might seem surprising. Many people seem to have two intuitions that work against the idea that different social networks can propel different ideas to the top: that quality always wins out in the end and that people will choose the best option regardless of other people's opinions. It turns out, however, that both of these intuitions are wrong.
- Results of a study done by sociologist Duncan Watts and his collaborators, Matthew Salganik and Peter Dodds, in the mid-2000s suggest that, at the very least, structures of social networks make a difference as to the quality of the information that you draw from those networks. In particular, they demonstrate that popularity on a network matters: Small differences in the popularity of an idea can make a big difference with respect to the spread of that idea, and that difference can't be explained through appeal to the intrinsic quality of the idea itself.
- This suggests that the structure of a social network is not irrelevant when analyzing the ideas that become popular on that social network. Furthermore, we can also use different measures for distinguishing between

good and bad ideas. Given our interest in knowledge, one obvious measure that we could use to distinguish between ideas would be in terms of how true or accurate they are.

- Taken altogether, this means that popularity on a social network is not a guarantee of how true or accurate that idea is. Different social network structures, in particular, can differ with respect to how well they do in making the true ideas popular and ensuring that the popular ideas in the social network are also true and accurate. In other words, if you're lucky enough to be in a social network in which the successful ideas are also true, then you'll end up endorsing many true ideas, but if you're unlucky and find yourself in a social network in which the successful ideas aren't all that reliably accurate, then you'll end up with a lot of misinformation.

- From the components of Dunbar's number, weak ties, and network effects, two big conclusions can be drawn about how we actually gain information from testimony.
  - The social groups in which we're embedded are often too large and complex for us to keep track of the individual reliability of each of the members of those social groups. The "too large" part of this claim is derived from the discussion of Dunbar's number. Even if you focus on the most widely cited number—150—that's too large for you to keep track, in a very systematic or reliable way, of how accurate each of those 150 people is with respect to all of the information he or she passes along. The "too complex" part of the claim comes from the importance of weak ties. It's hard enough to keep track of your casual acquaintances. But it's not possible for you to keep track of the casual acquaintances of your casual acquaintances' casual acquaintances—though it's these acquaintances that explain your weak ties. And given the importance of weak ties, these are the acquaintances that make a difference with respect to the sort of information that your social network gives you.
  - The structure of our social networks can contribute to our acquiring reliably accurate information just because of our embeddedness in those networks. The first reason is due, again, to Christakis and Fowler's discussion of weak social ties. According to them, a large amount of your information will be from people with whom you have no direct connection, not even casually. If you're lucky enough, however, to have a network such that even those with whom you're only connected through weak ties are reliably accurate informants, then you'll gain a lot of reliably accurate information, just by being embedded in your social

network. The second reason is due to the power of social networks to influence the ideas that we accept—which we do so in large measure because they are popular rather than because they are "good" or true.

Suppose there was a social network in which the popular ideas were true, and vice versa. If you were lucky enough to be embedded in such a network, you would end up believing a lot of reliably accurate information. It might be the case that you're believing the information because it's popular in your network, and not because it's true, or because you know it's true. But for our purposes right now, that's fine. The point would still remain: The structure of your social network would contribute to your reliably acquiring accurate information.

Our new theory is a form of externalism. It doesn't require that we consciously screen out unreliable sources from reliable ones, but it's a social externalism in that the screening process is completely off-loaded from the individual to the social network or networks in which he or she is embedded.

If the components of Dunbar's number, weak ties, and network effects are taken together, they suggest that a social network can serve to ensure that the information you receive from your network is reliably accurate. That's the foundation of social externalism.

READINGS

Carter, Clark, Kallestrup, Palermos, and Pritchard, eds., *Socially Extended Epistemology*.

Huebner, *Macrocognition*.

QUIZ

1 TRUE OR FALSE
Dunbar's numbers were predicted by computing the ratio of the volume of the neocortex, the part of the brain responsible for tracking complex information, to the total volume of the brain.

2 TRUE OR FALSE
In a study of word-of-mouth referrals of piano teachers reported in Nicholas Christakis and James Fowler's *Connected*, more than a third of referrals came from friends of friends of friends of the teachers' clients.

3 The spread of information across society depends on which of the following?

- **a** Many people having a lot of friends
- **b** A few people who are extremely social
- **c** Weak ties—connections with people who are a few degrees removed

4 TRUE OR FALSE
According to the work of Duncan Watts and his colleagues, social influence is irrelevant to the success of, for example, a song; only the quality of the song ultimately matters.

Answer key can be found on page 207.

LECTURE 20

The Reliability of Scientific Testimony

We shouldn't give up on the idea that relying on testimony can only produce knowledge if there is a way to ensure the testimony we're relying on is accurate. Instead, we should give up on the idea that the process ensuring that the testimony is accurate has to be located in the person relying on the testimony. And we should adopt a view called social externalism, introduced in the previous lecture. Socially distributed cognitive processes support the idea that social externalism is a genuine alternative both to traditional externalism and to Hume's non-presumptive theory of testimony.

SOCIAL EXTERNALISM

- Externalist theories of knowledge suggest that what makes a true belief count as an instance of knowledge is that the belief was formed by a reliably accurate cognitive process. So, what makes the position of social externalism different must have something to do with how it incorporates a social element into the theory of externalism. Thus, we can formulate a theory of social externalism by adding a notion of socially distributed cognitive processes to the theory of externalism that we already have.

- Simply put, a socially distributed cognitive process is a cognitive process carried out not within the mind or brain of an individual person, but extended across a social network.

- In a previous lecture, we were hesitant to endorse the view that the mind extends beyond the individual person. But now we're suggesting that we should appeal to socially distributed cognitive processes. How are these two positions compatible?

- The difference between the extended mind hypothesis and the current proposal is that extending the mind beyond the boundaries of the person does violence to the concept of mind, while adopting socially distributed cognitive processes does not do violence to the concept of cognition or the cognitive.

Mathematician Charles Babbage wrote that the use of inaccurate mathematical tables lost the British government "between two and three million" pounds sterling. Had the government invested just a fraction of that money on scientific research, Babbage suggested, such errors could easily have been avoided.

SOCIALLY DISTRIBUTED COGNITIVE PROCESSES

- While idea of socially distributed cognitive processes or systems has only really become a topic of research in the last few decades, the idea of such systems is more than 100 years old.

- Mathematician Charles Babbage was the first thinker to analyze socially distributed cognitive systems. In a chapter entitled "On the Division of Mental Labour" in his 1832 work *On the Economy of Machinery and Manufactures*, Babbage describes such a system in detail. In fact, it was Babbage's discussion of socially distributed cognitive systems that gave him the idea for development of tools for mechanical computation.

- In that chapter, Babbage describes a project organized by the French mathematician Gaspard de Prony to use a system of human "calculators" to compile a table of logarithms. The project was a huge success—extending over a number of years, with more than 100 workers assigned to different tasks—and eventually produced 18 volumes.

- The inaccuracy of mathematical tables was a significant impediment to scientific and technological progress—and a source of a great deal of financial insecurity—in the first half of the 19th century. This was the problem that de Prony solved by implementing a complex system of social organization designed specifically to achieve the most accurate tables of mathematical results known up to his time.

Here is Babbage's description of de Prony's design:

First Section.—The duty of this first section was to investigate, amongst the various analytical expressions which could be found for the same function, that which was most readily adapted to simple numerical calculation by many individuals employed at the same time. This section had little or nothing to do with the actual numerical work. When its labours were concluded, the formulae on the use of which it had decided, were delivered to the second section.

Second Section.—This section consisted of seven or eight persons of considerable acquaintance with mathematics: and their duty was to convert into numbers the formulae put into their hands by the first section,—an operation of great labour; and then to deliver out these formulae to the members of the third section, and receive from them the finished calculations. The members of this second section had certain means of verifying these calculations without the necessity of repeating, or even of examining the whole of the work done by the third section.

Third Section.—The members of this section, whose number varied from sixty to eighty, received certain numbers from the second section, and, using nothing more than simple addition and subtraction, they returned to that section the finished tables. It is remarkable that nine-tenths of this class had no knowledge of arithmetic beyond its two first rules which they were thus called upon to exercise, and that these persons were usually found more correct in their calculations, than those who possessed a more extensive knowledge of the subject.

Charles Babbage was one of the earliest researchers to explore the design of computers.

- Babbage saw de Prony's design as a model for future large-scale scientific investigations and called the tables that de Prony's workers achieved "one of the most stupendous monuments of arithmetical calculation which the world has yet produced."
- Babbage recognized that one aspect of de Prony's achievement rendered it almost unbelievable. As Babbage put it, the possibility of having achieved such a task employing only workers equipped with the most rudimentary arithmetical skills "may appear to non-mathematical readers to be rather too large a postulate," one shrouded in "apparent mystery."

A POTENTIAL OBJECTION TO SOCIAL EXTERNALISM

- In the previous lecture, we argued for social externalism because we suggested that socially distributed cognitive processes are potentially much better than lone individuals at filtering out unreliable information sources. Why, then, couldn't we require that individual people at least choose which socially distributed cognitive processes they are going to rely on based on how reliably accurate those socially distributed cognitive processes are?
- In effect, if we make this argument, we're suggesting that we should revive the Humean inferentialist theory, but now—instead of applying it to assessing individual informants—we should apply it to the socially distributed cognitive processes we rely on.
- Here are two reasons why this new application of Hume's inferentialist theory won't work.
  - We humans are not very good at source monitoring, which means we're not very good at remembering what the sources of our beliefs are. Applied to the socially distributed cognitive processes, this means we won't remember which of our beliefs we owe to particular cognitive processes, so we won't be able to correctly assess whether a cognitive process is reliably accurate.
  - Even the participants in socially distributed cognitive processes often don't know if they're successful, or why. The reason for this is that the aspects of socially distributed cognitive processes that contribute to their success are not always obvious—even to the participants in those processes. And if the people intimately involved in those processes themselves often don't know why those processes work, why would we expect outsiders to know?

In 1986, cognitive science researcher Naomi Miyake demonstrated that even a team composed only of two people, both of whom feel they are participating in a joint activity they each fully understand, can have very different interpretations of a relatively simple activity. Imagine how much greater the interpretations can diverge as the socially distributed processes become larger and more complex.

CONTEMPORARY SCIENCE

- Contemporary science provides us with examples of socially distributed cognitive processes. In fact, it would be a surprise if scientific researchers didn't take advantage of such cognitive processes. The challenges facing contemporary scientists and the advantages offered by large-scale collaborations to amplify the capacities and expertise of their individual members make implementing socially distributed cognitive systems in scientific research a natural fit.

- In fact, this is what Babbage predicted close to 200 years ago. He argued that the proof of concept in the industrialization of manufacturing would soon spread to science, writing that "[t]he progress of knowledge convinced the world that the system of the division of labour and of cooperation was as applicable to science, as it had been found available for the improvement of manufactures."

- For further evidence of the role of socially distributed cognitive processes in scientific research, we can also look to the work of John Seely Brown and Paul Duguid. In their highly influential 2000 book *The Social Life of Information*, they suggest that the complexity of knowledge creation, sharing, and deployment tasks involved in contemporary life favors the creation of socially distributed cognitive systems over the deployment of lone thinkers.

We think of big science—like big data—as a 21st century phenomenon. But Alvin Weinberg, a prominent director of Oak Ridge National Laboratory, used the term "big science" in an article published in *Science* magazine in 1961.

- So, if we consider the history of science as it progresses from Babbage's day to ours, individual scientists have gradually been squeezed out of the picture in favor of the socially distributed cognitive systems that now account for so much of scientific activity. In fact, there are entire cities—such as the United States' Oak Ridge National Laboratory, Japan's Tsukuba Academic City, and the Soviet Union's Akademgorodok—that arose in the 20th century to support scientific research. The analogy for the 21st century might be the "campuses" of Microsoft, Facebook, or Google, which are basically cities unto themselves.

- The evolution of socially distributed cognitive systems leads to increasingly more finely drawn divisions of labor within the task, thus decreasing the cognitive demands on each of the individuals within those systems and taking advantage of the features of the cognitive system itself to create expertise.

- In contemporary large-scale science, tasks have been portioned out among group members so as to reduce the complexity and computational demands required for the completion of those tasks. The widespread use of computers and other technological aids adds another dimension to the extent to which the results of scientific research can no longer be traced back to one solitary researcher—or even a small group of researchers.

- Some of the clearest examples of science as a socially distributed cognitive process are the massive collaborative efforts involved in some of the biggest achievements of contemporary science, such as the sequencing of the human genome or the confirmation of the existence of the Higgs boson. These efforts result in articles listing many collaborators—sometimes even hundreds or more.

- Contemporary scientific research meets our definition of a socially distributed cognitive process. The example of science provides strong support for the social externalist explanation for how testimony provides evidence for knowledge.

READINGS

Kitcher, *The Advancement of Science*.

Wagenknecht, *A Social Epistemology of Research Groups*.

QUIZ

1 TRUE OR FALSE
Without the mathematical expertise of all of the "calculators" involved in the project, Gaspard de Prony would never have been able to produce the volumes of mathematical tables that so impressed Charles Babbage.

2 TRUE OR FALSE
According to Edwin Hutchins's study of the sea-and-anchor detail, one officer kept constant track of all of the tasks performed by the team involved in the position-fixing task.

3 Which of the following is an example of socially distributed cognition?

- **a** Gaspard de Prony's production process for calculating mathematical tables
- **b** The sea-and-anchor detail described by Edwin Hutchins
- **c** The high-energy particle physics labs described by Karin Knorr Cetina
- **d** All of the above
- **e** None of the above

4 Which of the following explains why we can't reformulate David Hume's inferentialist theory of testimony to apply to socially distributed cognition?

- **a** Our unreliability at source monitoring
- **b** Naomi Miyake's research that even very small, two-person groups can have vastly divergent understandings of their joint activities
- **c** Donald Norman's observation that socially distributed cognitive systems evolve in ways that people fail to recognize
- **d** All of the above
- **e** None of the above

Answer key can be found on page 207.

LECTURE 21

Testimony in the Media

If social externalism is the correct theory of testimonial knowledge, then we have knowledge on the basis of the information we receive from socially distributed processes only if relying on those processes provides us with reliably accurate information. Do we, as theorists attempting to test whether social externalism offers a good theory of testimonial knowledge, have good reason to think there are any socially distributed processes that provide us with reliably accurate information?

FACT-CHECKING ORGANIZATIONS

- There are a small number of news organizations that devote significant resources to establishing a culture of fact-checking. That includes having a group of employees whose job is dedicated to fact-checking and who are independent from the editors and writers, who often have a vested interest in seeing a story go to press.

- Not only are there very few news organizations nowadays that have a strong culture of in-house fact-checking, but if you look at the history of news organizations, the practice of having journalistic fact-checkers has existed for less than a century.

- News organizations that have such strong commitments to fact-checking—such as *The New Yorker* and *Der Spiegel*—are plausible examples of socially distributed cognitive processes that reliably convey accurate information. Unfortunately, however, such organizations are not the norm. In addition, the number of publications to invest in a full-time fact-checking department has decreased significantly.

The German news magazine *Der Spiegel* has the largest fact-checking department of any publication in the world.

The process involved in fact-checking is quite rigorous. Fact-checkers for *The New Yorker* are so obsessive that they even fact-check the cartoons and cover art, as well as the fiction and poetry.

- The relative scarcity as well as the decline of cultures of fact-checking at media institutions means that we can't appeal to the practice of in-house fact-checking to make a general claim about the reliable accuracy of media organizations. What should we say about the media more broadly? Is it possible to be optimistic about the sort of information that you can expect to receive when you consider the media landscape as a whole?

- Only someone who hasn't been paying attention to developments over the last decade or more would be that optimistic. However, some have pointed to the rise in popularity of external fact-checking media organizations as a strategy for improving the overall accuracy of media offerings. Might those sorts of external checks on other media organizations at least help improve overall accuracy?

- We have good reason to be skeptical about putting too much faith in external fact-checking organizations. The problem is that once we rely on an external fact-checking organization to weigh in on the accuracy of a particular story, we also need to rely on individual consumers to consult those fact-checking organizations alongside the original news source they used to get their information in the first place.

- But there are multiple fact-checking organizations, and they often disagree about how to evaluate any particular story. In other words, by bringing in external fact-checking organizations, we are once again putting the responsibility on individuals to be able to assess information, gauge the relative strength of individual experts' opinions, and then arrive at a reliably accurate assessment of the totality of the information on balance.

- The CIA trains its analysts over the course of years to do that sort of assessment of source data. Trusting untrained people to do that same work is not a promising solution to the problem of unreliable source material.

- Added to this is the further problem of source monitoring failures—the problem that we're just not very good at remembering where we acquired a particular piece of information. For external fact-checking to even have a chance of working, people would have to read an individual story, stop, fact-check it using external resources, and then decide whether to believe it before going on to read the next individual story. And that's just not a plausible scenario.

- If that's the situation, the more plausible answer is probably to despair that we can acquire much knowledge from media sources at all.

Some of the publications that adhere to norms of rigorous fact-checking include *The New Yorker*, *Vanity Fair*, *Esquire*, *Wired*, and *Popular Mechanics*.

MULTIPLE MARKETS

- Perhaps an alternative solution to external fact-checking organizations would be to simply let the market weed out the bad actors.

- People think that what markets do is achieve the best-possible result at the best-possible price, but that is not exactly true. Even when they work well, markets achieve an optimal balance of the quality of the result and the cost of achieving that result. But even that isn't exactly right, because it doesn't tell us what aspects of a result count toward its quality.

- A car, for example, has different measures of quality, including its fuel efficiency, its maximal speed, its peak acceleration, and the attractiveness of its exterior and interior design. How do we determine which of those features is going to count toward its quality?

- The short answer is that we can't. But one of the advantages of market thinking is that we don't have to. In some sense, the market will do that for us: Whichever car sells best is the one that had the best combination of price and features.

- But when we say "best," all that means is popular with the consumers in that car-buying market. And here's what that means: When you see what features the market rewards, you learn something about what sorts of features car buyers wanted, rather than gaining information about some sort of objective standard for assessing which car is best, all things considered.

- And actually, as the case of the car indicates, there won't be one combination of qualities and price that will uniquely appeal to consumers. Instead, there will be a range of different combinations of qualities and price to attract different subsets of the car market. So, the car market is segmented into different categories that appeal to those different types of consumers.

- We can compare a market for ideas to the market for cars, and what we see is that—in the case of science, at least—we don't have the confusing range of qualities against which to measure the worth of an idea. So, unlike with cars, we don't have the potential sources of quality that we have to weigh against each other.

- In the case of science, the measure of the quality of an idea is whether it's true or likely to be true. That is the one measure of quality against which we have to weigh the worth of an idea. That does seem like an advantage for science and one reason why we could say that in science the market is selecting for the truth of the ideas under consideration.

- A complicating factor is that scientists and other researchers—including researchers in the humanities and social sciences—are actually participating in multiple markets at the same time. One sort of market is monetary, in which a researcher might be able to monetize his or her ideas and measure his or her success by the amount of money someone is willing to pay for him or her to pursue those ideas.

- At the same time, researchers also participate in what economists call reputation markets, in which the measure of the success of an idea isn't financial but reputational. It's the quality of the reputation that the idea and the originator of that idea have among other practitioners in the field.

- Furthermore, there can be other reputation markets in addition to the ones limited to individual academic fields. For example, certain high-profile researchers can enjoy a reputation in society more generally.

- The complication introduced by the fact that there are multiple markets is that there isn't a strong connection between the monetary and reputational markets. It is often the case that someone's ideas can be quite successful in reputation markets within his or her field without those ideas leading to corresponding financial success. And academic specialists will point out cases of their colleagues who have achieved financial success for their ideas without being able to achieve reputational success within their own field for those same ideas.

- Unfortunately, once you consider what the market is aimed at maximizing, it does not seem that a focus on the media marketplace would give us any reason to be optimistic that market mechanisms will make media organizations better sources of reliably accurate information.

- The reason for this is that what the media market is aimed at optimizing isn't reliably accurate information. It

aims at optimizing what it rewards—either with money or reputation. And it's geared toward rewarding whatever content can capture the attention of the greatest number of consumers.

- What this means for the news media is that their goal is to maximize attention—because by maximizing attention, they can deliver consumers and data about those consumers to advertisers and other corporate customers.

- This would not pose a problem if the way to maximize the attention of consumers was to provide them with a wealth of reliably accurate information. But that's not necessarily what consumers want.

In his 2016 book review "They've Got You, Wherever You Are," published in *The New York Review of Books*, Jacob Weisberg describes the problems created by the current media market using the example of Facebook.

His point is that making the media landscape more impoverished when it comes to sources of high-quality, reliable news may not be the intention of Facebook or any of the other big players in media markets. However, it is likely to be an unintended consequence, given that the market doesn't seem to be geared toward rewarding that sort of high-quality, reliable news.

SOURCES YOU RELY ON

- In order to assess whether you can acquire knowledge on the basis of your consumption of news media, we shouldn't look at all news sources. Instead, we should look at the news sources that you actually rely on.

- If you rely on *The New Yorker* and other sources that have strong journalistic traditions that emphasize accuracy, then you can acquire knowledge on the basis of relying on those sources.

> Your reliance on *The New Yorker* is a function of your position in a network of *New Yorker* readers and doesn't at all depend on your awareness of—or appreciation for—the fact that *The New Yorker* is a reliable source of accurate information.

- Externalism about knowledge has to do with whether you actually possess knowledge, not whether you or anybody else is aware that you possess that knowledge. So, our focus on how we would figure out whether you have knowledge is a little misleading.

- For the social externalist, in the case of your knowledge on the basis of reading a *New Yorker* article, what matters is whether relying on *New Yorker* articles is a good way to receive accurate information. It isn't important that you know that *The New Yorker* is reliable, or even that you're aware of that fact. Maybe you read it because you like the cartoons, but you also check out the articles. Or maybe you want to fit in with a group of people who are constantly discussing *New Yorker* articles.

- The fact that you're reading *The New Yorker* and that it's a source of reliably accurate information is enough for you to acquire knowledge on the basis of your reading. And this result is seen in other well-established cases in the social sciences.

Social externalism can be viewed as the most promising route to explain our knowledge on the basis of testimony.

READINGS

Allan, ed., *The Routledge Companion to News and Journalism*.

Goldman, *Knowledge in a Social World*.

QUIZ

1 Which of the following is not a reason why social externalism is a form of externalism?

a You don't need to be aware of the fact that you are relying on socially distributed processes for your information.

b Whether you actually acquire knowledge on the basis of your reliance on those processes depends on whether those processes provide you with reliably accurate information.

c You can acquire knowledge on the basis of relying on those socially distributed processes even if you're unaware that they are providing you with reliably accurate information.

d You have an unconscious awareness of the reliability of the processes on which you rely.

2 TRUE OR FALSE
In-house fact-checking is an institution that has existed since the first newspapers and magazines.

3 TRUE OR FALSE
The media marketplace selects for the most accurate, reliable information.

4 TRUE OR FALSE
If social externalism is true, then what matters for you to have knowledge isn't that all media sources are reliably accurate but rather that the media sources on which you rely are reliably accurate.

5 External fact-checking organizations are not as useful as in-house fact-checking organizations, from the point of view of promoting knowledgeable media consumers, because of which of the following?

a Consumers would then have to keep track of different external fact-checking organizations alongside the media that they consume.

b Accurately weighing the reliability of different sources is not something that we can expect of untrained media consumers.

c We are not good at source monitoring.

d All of the above.

e None of the above.

Answer key can be found on page 207.

LECTURE 22

Pragmatic and Moral Encroachment

We have been treating the discussion of knowledge as only revolving around questions of truth likelihood. On this way of looking at knowledge, what matters for knowledge is truth and accuracy. In particular, nothing else about the knower—including whether he or she has an interest in the particular question he or she is considering—is relevant to the question of whether he or she knows. This way of looking at knowledge as being independent of the interests of the knower is the standard. However, recently there have been a number of challenges to the claim that we can evaluate whether someone knows purely on the basis of the truth likelihood of his or her belief. According to these new challenges, whether someone knows something can also depend on his or her interests.

One challenge suggests that a person's practical interests can impact whether he or she has knowledge. The other challenge suggests that a person's moral concerns can impact whether he or she has knowledge. Though ultimately neither of these challenges will prove successful, both of them help illustrate not only the importance but also the limits of knowledge.

PRAGMATIC ENCROACHMENT

- A particular sort of argument for the claim that a person's practical interests can impact whether he or she counts as having knowledge has been advanced by philosophers Jeremy Fantl and Matthew McGrath. Other notable philosophers who have argued for related conclusions include John Hawthorne and Jason Stanley.
- Fantl and McGrath's argument goes under the heading of pragmatic encroachment, because it involves the conclusion that pragmatic, or practical concerns, influence or encroach upon questions of whether or not someone has knowledge.
- The structure of Fantl and McGrath's argument is very simple. Essentially, it rests on two planks: fallibilism and a practical condition on knowledge.
  - Fallibilism is the claim that it is possible for someone to know something without having certain knowledge. Given the fact that we've adopted fallibilism as a working assumption for the bulk of these lectures—ever since we rejected Descartes's extreme version of internalism—we'll continue with the assumption that fallibilism is correct.
  - According to the practical condition on knowledge, if you know some fact, then you are rational to act as if that fact is true.
- When you combine fallibilism and the practical condition, you are forced to accept that whether or not someone has knowledge depends in part on his or her practical interests.

- It's very plausible to think that whether it is rational to act as if some fact is true does vary according to your practical interests. The more significant the outcomes of a particular act are for you, the greater evidence you need in order to have reason to act.

- But then, if that's true, that would imply—because of the practical condition—that the evidence required for knowledge would also vary according to your practical interests.

- That, however, just is the claim endorsed under the heading of pragmatic encroachment. That is because what this conclusion means is that whether or not you have knowledge can depend on the level of importance a fact has for you. The more important the fact is for you, the more evidence you would need in order to know that fact. For trivial facts, though, you could have knowledge with far less evidence.

- Why should we accept the claim that whether it is rational for you to act as if a certain fact is true depends on your practical interests?

- Suppose you believe that whenever you eat peanuts you get headaches—not debilitating headaches, but annoying ones. If that's the case, then you might want to look out on menus for a disclaimer stating that the restaurant does not use peanuts or peanut products in preparing its food. Once you see that disclaimer, you'll feel comfortable ordering any dish on the menu.

- Contrast that with a situation in which you have a severe peanut allergy, one so severe that you immediately go into anaphylactic shock when you come into contact with even the smallest amount of peanut or peanut oil. If that's the case, then a mere disclaimer on a menu might not be enough for you to feel comfortable eating in that restaurant. Perhaps in that case you would need to speak to the restaurant manager, emphasize the severity of your condition, and then double-check to be sure the restaurant really uses no peanut products in their kitchen.

- In the first case, we might say that you would be irrational to continue to ask about the use of peanuts in the kitchen after reading a clear disclaimer on the menu. In the second case, however, nobody would consider you to be irrational if you made absolutely sure that the kitchen didn't use peanuts or peanut products.

- This seems plausible. And if it's correct, then it would suggest that whether or not your behavior counts as rational does depend on your practical interests.

The more interest you have in the outcome of a particular question, the more rational it is for you to invest additional time in acquiring more evidence related to the outcome of that question.

- So, if we're going to block Fantl and McGrath's argument that practical interests can influence whether or not you have knowledge, we'll have to block the argument by focusing on the practical condition—the claim that if you have knowledge of some fact, then it's rational to act as if that fact is true.

- When the practical condition is stated like that, it might seem really obvious. However, there are some good reasons to be skeptical about the practical condition.
  - It seems very plausible that many nonhuman animals have knowledge. When a hound tracks a fox to a particular tree, the hound can know that the fox is at that tree. At the same time, it is not plausible that most of the nonhuman animals that have knowledge are rational; in other words, the notion of rationality doesn't really apply to those animals. Rather, it seems plausible to reserve the notion of rationality for creatures who can reflect on their reasons for behaving the way they do. To say that someone is behaving rationally is at least in part to say that, if the person reflects on his or her behavior, he or she would be able to explain that behavior to him- or herself on the basis of his or her reasons for acting. It's no disrespect to the extraordinary cognitive abilities of nonhuman animals like dogs to suggest that they are incapable of such reflection. If that's right, though, then such animals provide examples of creatures who have knowledge but do not act rationally. And if that's right, then the practical condition is false.
  - According to the practical condition, whenever you know something, then that is enough for you to be rational in acting as if it's true. The problem is that this isn't true. Sometimes you need more than mere knowledge to have reason to act.

- It seems plausible that there are certain scenarios in which action requires more than mere knowledge. Sometimes, you have to go above and

beyond, acquiring evidence greater than what is required for knowledge.

- Think again about the discussion of the menu disclaimer that the kitchen at the restaurant where you're eating is peanut-free. If the restaurant goes to the trouble of indicating to you that it doesn't use any peanuts or peanut products in the kitchen, then you can know that the restaurant's kitchen is peanut-free.

- However, if you have a deadly peanut allergy, it's totally rational for you to speak to the manager to make absolutely sure that there are no peanuts or peanut products used in the kitchen. In fact, to go a step further, you shouldn't eat at that restaurant until you double-check that they really don't use any peanuts or peanut products in their kitchen.

- If you're willing to go that far, too, but you accept that the disclaimer in the menu is enough evidence for you to know that there are no peanuts or peanut products used in the kitchen, then you also must reject the practical condition. That's because this would also be a case in which knowledge alone isn't sufficient to provide you with reasons to act.

MORAL ENCROACHMENT

- Fantl and McGrath argue that whether or not you have knowledge of some fact can depend in part on what practical significance that fact has for you. The more significant the fact, the more evidence you might need to have in order to know it. It's this that they call pragmatic encroachment, because your pragmatic interests encroach on, or influence, the amount of knowledge you can have.

- There is another form of encroachment that has begun to receive interest: the encroachment of moral concerns on evidence. In analogy to the notion of pragmatic encroachment, we could call this view moral encroachment. It's the idea that the amount of evidence you need in order to have sufficient evidence to believe something depends in part on the moral implications of that belief.

- To unpack the notion of moral encroachment, let's emphasize that the idea of moral encroachment involves the suggestion that beliefs can have moral implications at all. And this may be hard for some people to accept.

- To say that beliefs can have moral implications is at least in part to say that certain beliefs can be morally wrong. In other words, the idea is that over and above being able to assess a belief as being supported

or unsupported by the evidence, we can also assess at least certain beliefs based on their moral qualities.

- If we accept the phenomenon of moral encroachment, we can't limit appeals to that phenomenon to people we agree with.

- Suppose, for example, we're debating with someone who believes that his or her religious tradition forbids women to work outside of the home. Say we present strong evidence to this person that children raised by working mothers are as healthy, happy, intelligent, and well adjusted as children who are raised by stay-at-home mothers. The person with whom we're debating could refuse to accept that evidence as sufficient, citing the moral significance of his or her religious injunction against women working outside of the home. We could imagine analogous cases involving white supremacists, anti-Semites, and so on.

- The underlying point is that we have good reason to separate arguments about the strength of our evidence from arguments about the rightness of our causes wherever possible. Of course, humans are often guilty of rationalization and sloppy thinking, but this isn't a reason not to try our best to shield our truth-directed investigations from our moralizing.

After reviewing the arguments that considerations beyond simply the likelihood that our beliefs are true might be relevant to our evidence, it doesn't seem that either of them is very promising.

READINGS

Fantl and McGrath, *Knowledge in an Uncertain World*.

Hawthorne, *Knowledge and Lotteries*.

QUIZ

1 Which of the following is not a presupposition of Jeremy Fantl and Matthew McGrath's argument that your knowledge depends on your practical interests?

- a It is possible to know something without being certain.
- b If you know some fact, then you are rational to act as if that fact is true.
- c It is possible to know something without knowing that you know.

2 TRUE OR FALSE

The fact that small children and nonhuman animals could have knowledge without rationality would be a counterexample to Jeremy Fantl and Matthew McGrath's practical condition.

3 The fact that your knowledge might depend on the moral implications of your belief is an example of which of the following?

- a Political correctness
- b Pragmatic encroachment
- c Moral encroachment

4 TRUE OR FALSE

According to Rima Basu and Mark Schroeder, there is something wrong with your apologizing for having a certain belief despite the fact that you believe yourself to have epistemically unimpeachable evidence.

Answer key can be found on page 207.

LECTURE 23

Radical Skepticism: The Brain in a Vat

Imagine that one night after you've fallen asleep, an evil scientist breaks into your bedroom, gives you an anesthetic, and spirits you away to his secret lab. Once there, he removes your brain, keeping it functioning by placing it in a specially designed vat of nutrient solution. He runs electrodes to your brain from a bank of supercomputers, feeding your brain impulses identical to those your brain would receive from your body. When you wake up, everything will seem the same to you. You'll go through your morning routine as you always do, feeling the water on your face when you wash up. What you won't know, however, is that you're no longer in that bedroom. You no longer have a face to wash. Instead, you're now a brain floating in a vat of nutrient solution being fooled into thinking that you're still waking up in your bedroom and going about your morning routine.

Skepticism

The word "skepticism" derives from the ancient Greek word *skeptikos*, which means "questioning" or "doubting."

Many skeptical arguments are natural, everyday affairs—such as when you question someone's evidence for a particular opinion about sports, or art, or politics. Philosophical skeptical arguments are related to those everyday forms of skepticism, but philosophical skeptical arguments tend to be more general. They attack a whole type of evidence, such as perceptual evidence, or inductive evidence, or the evidence of testimony.

Generally, such arguments suggest that the evidence in question isn't good enough to give us knowledge. For example, skepticism about sense perception suggests that the senses fool us into thinking they are providing us with evidence about the world around us. In reality, the skeptic suggests, the senses offer us only a false promise. The evidence that the senses provide is in fact not sufficient to support knowledge of the world that sense experience seems to describe.

Other types of skepticism work similarly. Whether you're dealing with skepticism about the existence of other minds or the existence of God, the nature of skeptical arguments is to question whether the evidence available is sufficient to support knowledge or belief.

THE BRAIN-IN-A-VAT SCENARIO

- Let's turn this brain-in-a-vat scenario into a skeptical argument, the conclusion of which will be that you can't have any knowledge of the world that you think you perceive.

- Let's consider a very obvious claim about the world that you believe to know through perception. Take, for example, the claim that you have hands.
  1. If you know that you have hands, then you know that you're not a (handless) brain in a vat.
  2. You don't know that you're not a brain in a vat.
  3. Therefore, you don't know that you have hands.

- But the claim that you have hands was chosen because it was an obvious claim about the world you think you know through perception. So, if the argument leads you to conclude that you don't even know *that* claim, then the argument would seem to be applicable to pretty much any claim about the world that you take yourself to perceive.

- At least initially, the argument seems pretty strong. It has the structure of a valid argument. Furthermore, the two steps of the argument seem unimpeachable.

- The first step of the argument seems plausible. Certainly, if you know that you have hands, then among other things you know that you're not some handless brain floating in a vat of nutrient solution.

- The second step of the argument also seems pretty strong. However the world seems to you right now, it would seem exactly the same even if you were merely a disembodied brain floating in a vat of nutrient solution and being fed experiences by an evil scientist. For this reason, you might think that you can't rule out the possibility you actually are a disembodied brain floating in a vat. And if you can't rule out that possibility, then surely you can't know that you're not in fact such a disembodied brain.

Two of the most famous discussions in the history of philosophy revolve around skepticism: Descartes's discussion of the evil demon and skepticism about our knowledge of the external world, and Hume's skepticism about the possibility of using induction to gain knowledge of relations of cause and effect.

- But if the argument is valid and if each of the steps of the argument is plausible, then we must accept the conclusion. So, it would seem that we must accept skepticism about the world as perceived through our senses.
- Or must we? If we're going to resist the external world skepticism prompted by the brain-in-a-vat argument, we must either reject one or both of the two steps or we must demonstrate that, despite appearances, the argument isn't in fact valid after all.
- Contemporary philosophers have pursued all three of these strategies, and each has something to teach us.

THE EXTERNALIST RESPONSE

- Externalists argue that facts about your situation—facts of which you are not aware—make a difference with respect to your evidence. So, two people who are the same with respect to the way the world seems to them can in fact have different levels of evidence.
- If this is right, then the externalist can reject the skeptical argument by rejecting the second step. You do in fact know that you're not a brain in a vat, assuming in fact that you really are a normally embodied human in your standard environment, perceiving the world through your bodily sense organs.
- One of the problems with the externalist response to the brain-in-a-vat argument is that if it is correct, it can seem mysterious why skepticism ever seemed so challenging in the first place. In order to have an explanation for this, we need to look at two other responses to the skeptic.

THE RELEVANT ALTERNATIVES VIEW

- The first of the two remaining responses rejects the first premise of the skeptical argument. It might seem outrageous to reject this premise. Isn't it just logic?

- A number of philosophers, foremost among them Robert Nozick and Fred Dretske, have questioned whether rejecting this premise is actually as outrageous as it might first appear. Nozick and Dretske reject the principle that underlies the first premise of the skeptical argument, called the closure of knowledge under known entailment.

Suppose you know the proposition that *p*. And suppose you further know that if *p* is true, then *q* is true—that is, that *p* entails *q*. It's plausible to suppose that you then also know the proposition that *q* is true.

For example, suppose you know that Georgina went to Caltech. And suppose you further know that if Georgina went to Caltech, then she has certainly mastered basic arithmetic. Then, it's plausible to suppose that you also know that Georgina has certainly mastered basic arithmetic.

This is the principle that knowledge is closed under known entailment: If you know some claim and you know that that claim entails some further claim, then you know the further claim as well.

Although this claim seems extremely plausible, it's the claim that Nozick and Dretske want to reject.

- They suggest that what's more plausible than the closure of knowledge under known entailment is this: When you have knowledge, that knowledge lets you rule out relevant alternatives that are incompatible with what you know. That's why the Nozick-Dretske view can plausibly be called a relevant alternatives view.

- They don't think that skeptical scenarios, for example, ever count as relevant alternatives. So, even though you know you have hands, you don't know the claim "if you have hands, then you're not a brain in a vat"—because the possibility that you are a brain in a vat is not a relevant alternative to your actual situation.

- Nozick and Dretske reject the closure of knowledge under known entailment and replace it with the relevant alternatives view.

- If that's right, though, then you have reason to reject the skeptical argument, even if you sympathize with the claim that you don't know you're not a brain in a vat. You can reject the argument, because you think the radical skeptical scenario is not relevant to your real-world knowledge that you have hands, despite the fact that having hands isn't compatible with being a brain in a vat—and you know that it's not compatible.

- The relevant alternatives view gives you a way to concede to the skeptic that you have no way to refute radical skeptical claims but nevertheless to hang on the everyday knowledge that you know yourself to have.

CONTEXTUALISM

- For many philosophers, rejecting the closure of knowledge under entailment is a step too far. They just don't find it plausible that you could know some claim and know that that claim entails some other claim yet fail to know the entailed claim.

- For those philosophers, there is an alternate theory that employs intuitions similar to the ones Nozick and Dretske appeal to but that doesn't require that you reject the closure of knowledge under entailment.

- The new theory suggests that Nozick and Dretske were wrong to apply their intuition to a particular premise—the first premise—of the skeptical argument. Instead, the new theory says we should apply the operating intuition that Nozick and Dretske appeal to the skeptical argument as a whole.

- Consider a normal, everyday scenario in which you're not thinking about skepticism. Say you're thinking about purchasing gloves. In a scenario

like that, you'd obviously know you have hands—you're contemplating purchasing a new pair of gloves. In such a normal scenario, you also know, although merely unconsciously, that if you have hands, you're not a brain in a vat. This new theory says that in such an everyday scenario, you also know, also merely unconsciously, that you're not a brain in a vat.

- Now consider a scenario where you're arguing explicitly with a skeptic. In this scenario, you know—now consciously—that if you know you have hands, then you know you're not a brain in a vat. But now, because you're consciously considering the skeptical possibilities, you don't know that you're not a brain in a vat. So now, when you're arguing with the skeptic, you don't know that you have hands.

- In other words, on this new theory, what you know varies according to the situation in which you find yourself—your knowledge varies according to your context. For this reason, this view of knowledge is called contextualism.

- The advantage of contextualism is that it explains the intuitive pull of skepticism without forcing us to abandon the idea that we can use inference to extend the scope of our knowledge unrestrictedly.

- The contextualist distinguishes between high-stakes and low-stakes cases and suggests that what is different between the two is the level of evidence you need to have in order to have knowledge. In the high-stakes case, you need a higher level of evidence, while in the low-stakes case, you need a lower level of evidence.

- The views of contextualism and pragmatic encroachment—and the arguments used to support them—are parallel. But while pragmatic encroachment deals with the connection between knowledge, rationality, and action, contextualism deals with the connection between knowledge and your ability to assert that you have knowledge. The contextualist suggests that in high-stakes cases, you need much higher levels of evidence to be able to assert that you have knowledge.

- One reason the contextualist answer to the skeptic is not one we should adopt is that the same arguments that we made against pragmatic encroachment also speak against coherentism. It might be that despite the fact that you have knowledge, you need higher levels of evidence to assert your knowledge in high-stakes cases. But in the context of skepticism, perhaps the central reason for rejecting contextualism is that as an answer to the skeptic, the contextualist response is not a refutation—instead, it's a capitulation.

- Think about the contextualist explanation of what goes on when confronted with a skeptical argument. According to the contextualist, whenever the skeptic raises his or her skeptical worries, we no longer have the knowledge that we would in nonskeptical contexts. In other words, rather than showing how the skeptic is misguided, as far as the contextualist is concerned, the skeptic always wins every argument.

- If the goal is to explain where the skeptic has gone astray, the contextualist response is disastrous.

- Luckily, we can take advantage of the positive aspects of the contextualist explanation of the strength of skeptical arguments without conceding everything to the skeptic that the contextualist does.

- In order to explain the cases, we don't need to say that you lack knowledge in the high-stakes case. Instead, we can say that what you lack in the high-stakes case is the evidence sufficient for confidence, or assurance. So, the high-stakes case, on this alternate explanation, would be one where you have knowledge but lack the confidence or assurance to assert your knowledge. This is a way to repurpose the refutation of pragmatic encroachment as a strategy for rejecting contextualism as well.

- This alternate explanation also allows us to explain what is happening in skeptical cases without capitulating to the skeptic. When the skeptic raises his or her doubts, the context becomes one in which we are no longer capable of asserting with assurance that we have knowledge. However—unlike in the contextualist explanation—this alternate explanation does not require that we grant that we no longer have knowledge. Instead, we can say that we still have knowledge when arguing with the skeptic, even if we cannot confidently defend the knowledge that we possess.

READINGS

DeRose and Warfield, *Skepticism*.

Hazlett, *A Critical Introduction to Skepticism*.

QUIZ

1 Rejecting the brain-in-a-vat argument because you do know you're not a brain in a vat is characteristic of which of the following?

- a Externalism
- b Rejection of the closure of knowledge under known entailment
- c Contextualism

2 Rejecting the brain-in-a-vat argument because you could know that you have hands without knowing that you're not a brain in a vat is characteristic of which of the following?

- a Externalism
- b Rejection of the closure of knowledge under known entailment
- c Contextualism

3 Rejecting the brain-in-a-vat argument in everyday scenarios but accepting that the skeptic wins whenever he or she explicitly asserts the skeptical argument is characteristic of which of the following?

- a Externalism
- b Rejection of the closure of knowledge under known entailment
- c Contextualism

4 TRUE OR FALSE

One way to reject contextualism is to say that sometimes you can know something without having enough evidence to assert what you know.

Answer key can be found on page 207.

LECTURE 24

The Future of Epistemology

The sources of knowledge can roughly be divided into personal and social sources. Epistemology can contribute to the study of each of these types of sources, and the future contributions of epistemology in these areas can help us recognize that all of the major theories of knowledge found in this course—foundationalism, coherentism, and externalism—have a role to play in the future development of the field.

The lectures in this course have covered a lot of material, but there is a wealth of material that exists that the course could not cover. Two of the most active current areas of research in epistemology might spark your interest in exploring the field further after you conclude these lectures.

- The success of Timothy Williamson's book *Knowledge and Its Limits* points to a renewed interest in epistemology in using formal, logical, and mathematical methods to clarify problems. This subfield within epistemology is called formal epistemology.
- In 2007, philosophy professor Miranda Fricker published a book entitled *Epistemic Injustice: Power and the Ethics of Knowing*, in which she suggests the phenomenon that she terms epistemic injustice comes in two forms: testimonial injustice and hermeneutical injustice.
  - Testimonial injustice harms speakers because it prevents them from being believed for reasons having to do with prejudice, rather than for valid reasons having to do with their qualities as honest or reliably accurate informants.
  - Hermeneutical injustice harms people because it prevents them from correctly interpreting or recognizing facts about themselves or about their situation.

PERSONAL SOURCES OF KNOWLEDGE

- Personal sources of knowledge include self-awareness, sense perception, rational reflection, and deductive and inductive inference.

- As noted in the discussion of each of those sources of knowledge in previous lectures, present research in philosophy benefits from extensive cross-pollination with the scientific fields investigating these phenomena. In fact, it is very unusual nowadays to find a philosopher working on epistemology and sense perception, for example, who isn't conversant with current work on the neuroscience of vision. Similarly, the best work on the epistemology of memory is informed by the latest studies in the cognitive psychology and neurobiology of memory.

- For example, for more than a decade, philosopher Gilbert Harman and electrical engineer Sanjeev Kulkarni have been collaborating on joint research. Since the beginning of their collaboration, they have coauthored two books: *Reliable Reasoning: Induction and Statistical Learning Theory* and *An Elementary Introduction to Statistical Learning Theory.*

- Statistical learning theory is at the root of contemporary approaches to machine learning and is fundamental to the use of artificial intelligence in applications in image recognition, speech recognition, medical diagnostics, and finance. When Facebook tags your face in a friend's post or when you speak to your smart home speaker, you have statistical learning theory to thank for those applications.

- Harman's collaboration with Kulkarni is just one example of the ways in which epistemologists have directly or indirectly influenced the development of disciplines in related fields, from cognitive science to computer science.

SOCIAL SOURCES OF KNOWLEDGE

- In the study of personal sources of knowledge, epistemologists have a long tradition of contributing behind the scenes to the development of many of the cognitive and behavioral sciences. However, some of the most exciting work for epistemologists in the future will perhaps relate to the investigation of the social sources of knowledge.

- One case builds on the consideration of the media in a previous lecture. Current examinations of misinformation and so-called fake news tend to focus on the ways that individual media consumers can guard themselves against the effects of misinformation. Primarily, the suggestion is that media consumers need to become better at critical thinking skills.

- For example, a *Forbes* magazine article from February 2018 entitled "How to Maintain Critical Thinking in the Modern World of New Media" provides a list of tips for ways that readers can become more critically savvy. But if you remember the discussion of social psychology and testimony from a previous lecture, you'll know that attempts to improve people's ability to detect deception and inaccurate information just aren't very effective.

- The lesson from the discussion of testimony in a previous lecture was that there is no reason to think that a solution to the problem of accurate information transmission will be focused on individual consumers of that information. Instead, we ought to look at social and institutional remedies.

- The subfield concerned with these sorts of questions—social epistemology—is a growing research area in epistemology that is still in its infancy. One of the giants in the field is philosophy professor Alvin Goldman, who, in his book *Knowledge in a Social World*, argues for a position that he terms veritistic social epistemology. In fact, social externalism is a form of veritistic social epistemology.

> "Veritistic" comes from the Latin word for truth, *veritas*. Veritistic social epistemology suggests that we should evaluate social institutions and processes to see whether they reliably contribute to the acquisition of true beliefs by individuals.

- Although it was published in 1999, Goldman's book contains a wealth of discussions that are still relevant to our media landscape today. For example, Goldman discusses a proposal to use strategic tax policy to alleviate the effects of advertising influences on print media like newspapers and magazines.

- Let's modify that discussion for the contemporary media landscape. The goal is to increase the amount of reliably accurate information in the information marketplace.

- Currently, content purveyors like Google and Facebook profit from, to put it bluntly, selling their users to advertisers. However, traditional news sources still produce much of the content that Google and Facebook provide. In other words, the traditional news sources are producing the content, but Google and Facebook are receiving the advertising revenue. This is not a sustainable model.

- One potential solution would be to treat news content analogously to music content. Each time you listen to a song on Spotify or iTunes, the artist gets a percentage of the profit that Spotify or iTunes receives from you. Something similar could work in the case of news organizations: Each time a user of Facebook or Google reads a news story, the news organization that produced the story could receive a percentage of the advertising revenue that Facebook or Google earned on the story.

- One of the problems with this suggestion is that although it would provide increased financial security for traditional news media sources, it would create other problems. For example, it might increase the already-existing pressure on those sources to produce clickbait-style content in order to drive the new revenue stream from Google and Facebook. From a truth-oriented perspective, that would certainly not be welcome.

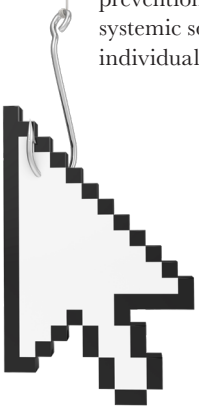

- These issues are challenging. Nevertheless, there are two advantages to the systems-oriented approach.
  - The alternative—focusing on making individuals more rational and reflective—isn't very promising. You can't fight human nature!
  - There are ways that systemic changes can produce positive truth-oriented effects. For example, the institution of cultures of fact-checking in American publications in the 1920s led to some more-reliable media institutions.

- The problem of the spread of misinformation is a systemic problem. And as with so many of the systemic challenges that we face—from global climate change to infectious disease prevention—we will need to find systemic solutions as opposed to individual approaches.

A FRAMEWORK FOR THINKING ABOUT VIEWS OF KNOWLEDGE

- We've seen two different ways to think about the structure of knowledge: the opposition between foundationalism and coherentism, and the clash between internalism and externalism. There is a way to reconcile some of these oppositions and take advantage of some of the best features of these views.

- To do this, let's borrow a set of ideas from philosophy professor Ernest Sosa. In a highly influential essay entitled "The Raft and the Pyramid," published in 1980, Sosa suggests that the best explanation of the structure of knowledge would combine elements from both foundationalism and coherentism. Basically, the structure that Sosa advocated involved having a foundational substructure supporting a coherentist superstructure.

- In recent decades, in a series of books and articles, Sosa has reworked that earlier picture. His new emphasis is on a distinction between two types of knowledge that he refers to as animal and reflective knowledge. Roughly, on this new picture, we should think of animal knowledge as foundationalist and externalist, while reflective knowledge is internalist and more coherentist.

- Sosa had the right idea but was thinking at the wrong scale. He thinks that the way to integrate coherentism, foundationalism, and externalism is by looking at a single individual thinker. That thinker has, according to Sosa, animal knowledge—his or her core, foundational knowledge. That type of knowledge is knowledge we should understand by appealing to externalism. Then, Sosa continues, the thinker has a higher level of knowledge that is validated by reflection. And Sosa thinks we should understand that type of knowledge at least in part using the resources of coherentism.

- In contrast, when it comes to knowledge, we should understand individual human knowledge largely on the model of externalism. Most human thinkers are not very reliable when it comes to rational reflection and shaping their beliefs into a truth-conducive, coherent whole.

- Remember, if we're talking about knowledge, it's not enough merely that you have coherent beliefs. Instead, the sort of coherence you achieve also has to lead to truth. That's a very high bar, and probably not many people can actually achieve that.

- Instead, what Sosa calls reflective knowledge happens at the level of groups, social processes, and institutions. The best scientific research groups and the most respected journalistic organizations are examples of how groups can achieve the sort of reflective processes that are conducive to truth.

- The role of internalist foundationalism and coherentism, then, will be in studying the kinds of processes that, when implemented at a systemic level, are likely to lead to truth.

- Of course, we shouldn't stop there. Here, too, empirical disciplines will have a role to play in testing whether implementing those processes actually does help with our truth-seeking goals and will also help us determine if there are other unforeseen costs that might outweigh the benefits of achieving those truth-seeking goals.

- What we're left with, then, is a sort of epistemology that is both externalist and foundationalist and coherentist. It is also an epistemology that sees its role in collaboration with the natural and social sciences. With all of the challenges to knowledge posed by contemporary life, it is a sort of epistemology that has a lot of work to do.

READINGS

Pritchard and Hendricks, eds., *New Waves in Epistemology*.

Steup, Turri, and Sosa, eds., *Contemporary Debates in Epistemology*.

QUIZ

1 The claim that if you know something, then you know that you know it is called what?

a Infallibilism
b Internalism
c Externalism
d The KK thesis

2 If you refuse to believe someone's testimony solely because of that person's membership in a disadvantaged group rather than because of any reasons pertaining to his or her honesty or reliable accuracy, that's an example of what?

a Formal epistemology
b Testimonial injustice
c Hermeneutical injustice

3 TRUE OR FALSE

Alvin Goldman's veritistic social epistemology suggests that we need to account for aspects of social information processes other than the truth accuracy of the information transmitted by those processes.

4 According to Ernest Sosa, what is the higher type of knowledge that should be distinguished from animal knowledge?

a Foundational knowledge
b Externalist knowledge
c Coherentist knowledge
d Reflective knowledge

Answer key can be found on page 207.

Answer Key

CLICK to navigate.

LECTURE 1
1 F 2 C 3 F 4 B 5 C

LECTURE 2
1 C 2 D 3 F 4 F 5 T

LECTURE 3
1 T 2 T 3 F 4 D

LECTURE 4
1 F 2 B 3 F 4 D

LECTURE 5
1 T 2 T 3 F 4 C

LECTURE 6
1 T 2 F 3 A 4 B 5 B

LECTURE 7
1 T 2 T 3 F 4 C 5 B

LECTURE 8
1 A 2 F 3 T 4 A 5 B

LECTURE 9
1 B 2 T 3 A 4 T 5 T 6 C

LECTURE 10
1 A 2 T 3 T 4 D

LECTURE 11
1 D 2 T 3 F 4 D 5 T

LECTURE 12
1 B 2 C 3 B 4 B 5 B 6 C

LECTURE 13
1 A 2 B 3 D 4 B 5 A

LECTURE 14
1 D 2 T 3 D 4 T

LECTURE 15
1 C 2 F 3 T 4 F 5 C

LECTURE 16
1 B 2 A 3 T 4 C

LECTURE 17
1 A 2 B 3 C 4 C

LECTURE 18
1 T 2 C 3 F 4 D

LECTURE 19
1 T 2 T 3 C 4 F

LECTURE 20
1 F 2 F 3 D 4 D

LECTURE 21
1 D 2 F 3 F 4 T 5 D

LECTURE 22
1 C 2 T 3 C 4 T

LECTURE 23
1 A 2 B 3 C 4 T

LECTURE 24
1 D 2 B 3 F 4 D

Bibliography

Adams, Frederick, and Kenneth Aizawa. *The Bounds of Cognition*. Blackwell Publishing, 2008.

Allan, Stuart, ed. *The Routledge Companion to News and Journalism*. Routledge, 2009.

Alston, William P. *A Realist Conception of Truth*. Cornell University Press, 1997.

Audi, Robert. *Epistemology: A Contemporary Introduction to the Theory of Knowledge*. Routledge, 2010.

Bengson, John, and Marc A. Moffett, eds. *Knowing How: Essays on Knowledge, Mind, and Action*. Oxford University Press, 2011.

Bernecker, Sven. *Memory: A Philosophical Study*. Oxford University Press, 2010.

Bernecker, Sven, and Kourken Michaelian, eds. *The Routledge Handbook of Philosophy of Memory*. Routledge, 2017.

BonJour, Laurence. *The Structure of Empirical Knowledge*. Harvard University Press, 1985.

BonJour, Laurence, and Ernest Sosa. *Epistemic Justification: Internalism vs. Externalism, Foundations vs. Virtues*. Blackwell, 2003.

Carter, J. Adam, Andy Clark, Jesper Kallestrup, S. Orestes Palermos, and Duncan Pritchard, eds. *Socially Extended Epistemology*. Oxford University Press, 2018.

Casullo, Albert. *A Priori Justification*. Oxford University Press, 2003.

Chisholm, Roderick. *Theory of Knowledge*. 2nd ed. Prentice-Hall, 1977.

Coady, C. A. J. *Testimony: A Philosophical Study*. Oxford University Press, 1992.

Conee, Earl, and Richard Feldman. *Evidentialism: Essays in Epistemology*. Clarendon Press, 2004.

DeRose, Keith, and Ted A. Warfield. *Skepticism: A Contemporary Reader*. Oxford University Press, 1999.

Fantl, Jeremy, and Matthew McGrath. *Knowledge in an Uncertain World*. Oxford University Press, 2009.

Gelfert, Axel. *A Critical Introduction to Testimony*. Bloomsbury Publishing, 2014.

Gertler, Brie. *Self-Knowledge*. Routledge, 2011.

Goldman, Alvin I. *Epistemology and Cognition*. Harvard University Press, 1986.

———. *Knowledge in a Social World*. Oxford University Press, 1999.

Harman, Gilbert. *Change in View*. MIT Press, 1986.

Hawthorne, John. *Knowledge and Lotteries*. Oxford University Press, 2004.

Hazlett, Allan. *A Critical Introduction to Skepticism*. Bloomsbury Publishing, 2014.

Huebner, Bryce. *Macrocognition: A Theory of Distributed Minds and Collective Intentionality*. Oxford University Press, 2013.

Hurlburt, Russell, and Eric Schwitzgebel. *Describing Inner Experience? Proponent Meets Skeptic*. MIT Press, 2011.

Jackson, Frank. *Perception: A Representative Theory*. Cambridge University Press, 1977.

Kitcher, Philip. *The Advancement of Science: Science without Legend, Objectivity without Illusion*. Oxford University Press, 1993.

Lackey, Jennifer, and Ernest Sosa, eds. *The Epistemology of Testimony*. Oxford University Press, 2006.

Lehrer, Keith. "Coherence, Truth, and Undefeated Justification." Chap. 7 in *Theory of Knowledge*, 2nd ed. Westview Press, 2000.

———. *Knowledge*. Clarendon Press, 1974.

Locke, Don. *Memory*. Macmillan, 1971.

Menary, Richard, ed. *The Extended Mind*. MIT Press, 2010.

Mercier, Hugo, and Dan Sperber. *The Enigma of Reason*. Harvard University Press, 2017.

Michaelian, Kourken, Dorothea Debus, and Denis Perrin, eds. *New Directions in the Philosophy of Memory*. Routledge, 2018.

Moser, Paul K., ed. *A Priori Knowledge*. Oxford University Press, 1987.

Musgrave, Alan. *Common Sense, Science and Scepticism: A Historical Introduction to the Theory of Knowledge*. Cambridge University Press, 1993.

Noe, Alva. *Action in Perception*. MIT Press, 2004.

Paul, L. A. *Transformative Experience*. Oxford University Press, 2015.

Pritchard, Duncan H., and Vincent Hendricks, eds. *New Waves in Epistemology*. Palgrave Macmillan, 2007.

Russell, Bertrand. *The Problems of Philosophy*. William and Norgate, 1912.

Shieber, Joseph. *Testimony: A Philosophical Introduction*. Routledge, 2015.

Skyrms, Brian. *Choice and Chance*. Wadsworth, 2000.

Stalker, Douglas F., ed. *Grue: The New Riddle of Induction*. Open Court, 1994.

Stanley, Jason. *Knowledge and Practical Interests*. Oxford University Press, 2005.

Steup, Matthias, John Turri, and Ernest Sosa, eds. *Contemporary Debates in Epistemology*. 2nd ed. John Wiley & Sons, 2013.

Stroud, Barry. *Hume*. Routledge & Keegan Paul, 1977.

Swinburne, Richard, ed. *The Justification of Induction*. Oxford University Press, 1974.

Wagenknecht, Susann. *A Social Epistemology of Research Groups*. Palgrave Macmillan, 2016.

Image Credits

5: sturti/E+/Getty Images; 12: drbimages/E+/Getty Images; 13: thegreekphotoholic/Getty Images; 22: PHOTOS.com>>/Getty Images Plus; 24: PeopleImages/Digital Vision/ Getty Images; 29: GettyImages-87572447.jpg; 37: Jupiterimages/PHOTOS.com>>/Getty Images Plus; 42: Jupiterimages/PHOTOS.com>>/Getty Images Plus; 46: MmeEmil/E+/Getty Images; 55: -AZ-/E+/Getty Images; 57: Pixelci/iStock/Getty Images Plus; 61: Dr_Microbe/iStock/Getty Images Plus; 75: StockPlanets/E+/Getty Images; 86: PeopleImages/E+/Getty Images; 105: Giulio Fornasar/iStock/Getty Images Plus; 127: BirdImages/E+/Getty Images; 130: VvoeVale/iStock/Getty Images Plus; 146: Wavebreakmedia/iStock/Getty Images Plus; 160: republica/iStock/Getty Images Plus; 167: Rawpixel/iStock/Getty Images Plus; 168: PHOTOS.com>>/Getty Images Plus; 176: Rawpixel/iStock/Getty Images Plus; 179: CASEZY/iStock/Getty Images Plus; 192: kickstand/E+/Getty Images; 203: Nerthuz/iStock/Getty Images Plus