

Philosophy of Science

Part I

Professor Jeffrey L. Kasser



THE TEACHING COMPANY ®

Jeffrey L. Kasser, Ph.D.

Teaching Assistant Professor, North Carolina State University

Jeff Kasser grew up in southern Georgia and in northwestern Florida. He received his B.A. from Rice University and his M.A. and Ph.D. from the University of Michigan (Ann Arbor). He enjoyed an unusually wide range of teaching opportunities as a graduate student, including teaching philosophy of science to Ph.D. students in Michigan's School of Nursing. Kasser was the first recipient of the John Dewey Award for Excellence in Undergraduate Education, given by the Department of Philosophy at Michigan. While completing his dissertation, he taught (briefly) at Wesleyan University. His first "real" job was at Colby College, where he taught 10 different courses, helped direct the Integrated Studies Program, and received the Charles Bassett Teaching Award in 2003.

Kasser's dissertation concerned Charles S. Peirce's conception of inquiry, and the classical pragmatism of Peirce and William James serves as the focus of much of his research. His essay "Peirce's Supposed Psychologism" won the 1998 essay prize of the Charles S. Peirce Society. He has also published essays on such topics as the ethics of belief and the nature and importance of truth. He is working (all too slowly!) on a number of projects at the intersection of epistemology, philosophy of science, and American pragmatism.

Kasser is married to another philosopher, Katie McShane, so he spends a good bit of time engaged in extracurricular argumentation. When he is not committing philosophy (and sometimes when he is), Kasser enjoys indulging his passion for jazz and blues. He would like to thank the many teachers and colleagues from whom he has learned about teaching philosophy, and he is especially grateful for the instruction in philosophy of science he has received from Baruch Brody, Richard Grandy, James Joyce, Larry Sklar, and Peter Railton. He has also benefited from discussing philosophy of science with Richard Schoonhoven, Daniel Cohen, John Carroll, and Doug Jesseph. His deepest gratitude, of course, goes to Katie McShane.

Table of Contents

Philosophy of Science

Part I

| | |
|---|----------|
| Professor Biography | i |
| Course Scope | 1 |
| Lecture One Science and Philosophy | 3 |
| Lecture Two Popper and the Problem of Demarcation | 3 |
| Lecture Three Further Thoughts on Demarcation | 9 |
| Lecture Four Einstein, Measurement, and Meaning | 12 |
| Lecture Five Classical Empiricism | 14 |
| Lecture Six Logical Positivism and Verifiability | 16 |
| Lecture Seven Logical Positivism, Science, and Meaning | 19 |
| Lecture Eight Holism | 22 |
| Lecture Nine Discovery and Justification | 25 |
| Lecture Ten Induction as Illegitimate | 28 |
| Lecture Eleven Some Solutions and a New Riddle | 31 |
| Lecture Twelve Instances and Consequences | 34 |
| Timeline | 37 |
| Glossary | Part II |
| Biographical Notes | Part III |
| Bibliography | Part III |

Philosophy of Science

Scope:

With luck, we'll have informed and articulate opinions about philosophy and about science by the end of this course. We can't be terribly clear and rigorous prior to beginning our investigation, so it's good that we don't need to be. All we need is some confidence that there is something about science special enough to make it worth philosophizing about and some confidence that philosophy will have something valuable to tell us about science. The first assumption needs little defense; most of us, most of the time, place a distinctive trust in science. This is evidenced by our attitudes toward technology and by such notions as who counts as an expert witness or commentator. Yet we're at least dimly aware that history shows that many scientific theories (indeed, almost all of them, at least by one standard of counting) have been shown to be mistaken. Though it takes little argument to show that science repays reflection, it takes more to show that philosophy provides the right tools for reflecting on science. Does science need some kind of philosophical grounding? It seems to be doing fairly well without much help from us. At the other extreme, one might well think that science occupies the entire realm of "fact," leaving philosophy with nothing but "values" to think about (such as ethical issues surrounding cloning). Though the place of philosophy in a broadly scientific worldview will be one theme of the course, I offer a preliminary argument in the first lecture for a position between these extremes.

Although plenty of good philosophy of science was done prior to the 20th century, nearly all of today's philosophy of science is carried out in terms of a vocabulary and problematic inherited from logical positivism (also known as logical empiricism). Thus, our course will be, in certain straightforward respects, historical; it's about the rise and (partial, at least) fall of logical empiricism. But we can't proceed purely historically, largely because logical positivism, like most interesting philosophical views, can't easily be understood without frequent pauses for critical assessment. Accordingly, we will work through two stories about the origins, doctrines, and criticisms of the logical empiricist project. The first centers on notions of meaning and evidence and leads from the positivists through the work of Thomas Kuhn to various kinds of social constructivism and postmodernism. The second story begins from the notion of explanation and culminates in versions of naturalism and scientific realism. I freely grant that the separation of these stories is somewhat artificial, but each tale stands tolerably well on its own, and it will prove helpful to look at similar issues from distinct but complementary angles. These narratives are sketched in more detail in what follows.

We begin, not with logical positivism, but with a closely related issue originating in the same place and time, namely, early-20th-century Vienna. Karl Popper's provocative solution to the problem of distinguishing science from pseudoscience, according to which good scientific theories are *not* those that are highly confirmed by observational evidence, provides this starting point. Popper was trying to capture the difference he thought he saw between the work of Albert Einstein, on the one hand, and that of such thinkers as Sigmund Freud, on the other. In this way, his problem also serves to introduce us to the heady cultural mix from which our story begins.

Working our way to the positivists' solution to this problem of demarcation will require us to confront profound issues, raised and explored by John Locke, George Berkeley, and David Hume but made newly urgent by Einstein, about how sensory experience might constitute, enrich, and constrain our conceptual resources. For the positivists, science exhausts the realm of fact-stating discourse; attempts to state extra-scientific facts amount to metaphysical discourse, which is not so much false as meaningless. We watch them struggle to reconcile their empiricism, the doctrine (roughly) that all our evidence for factual claims comes from sense experience, with the idea that scientific theories, with all their references to quarks and similarly unobservable entities, are meaningful and (sometimes) well supported.

Kuhn's historically driven approach to philosophy of science offers an importantly different picture of the enterprise. The logical empiricists took themselves to be explicating the "rational core" of science, which they assumed fit reasonably well with actual scientific practice. Kuhn held that actual scientific work is, in some important sense, much less rational than the positivists realized; it is driven less by data and more by scientists' attachment to their theories than was traditionally thought. Kuhn suggests that science can only be understood "warts and all," and he thereby faces his own fundamental tension: Can an understanding of what is intellectually special about science be reconciled with an understanding of actual scientific practice? Kuhn's successors in sociology and philosophy wrestle (very differently) with this problem.

The laudable empiricism of the positivists also makes it difficult for them to make sense of causation, scientific explanation, laws of nature, and scientific progress. Each of these notions depends on a kind of connection or structure that is not present in experience. The positivists' struggle with these notions provides the occasion for our second narrative, which proceeds through new developments in meaning and toward scientific realism, a view that seems as commonsensical as empiricism but stands in a deep (though perhaps not irresolvable) tension with the latter position. Realism (roughly) asserts that scientific theories can and sometimes do provide an accurate picture of reality, including unobservable reality. Whereas constructivists appeal to the theory-dependence of observation to show that we help constitute reality, realists argue from similar premises to the conclusion that we can track an independent reality. Many realists unabashedly use science to defend science, and we examine the legitimacy of this naturalistic argumentative strategy. A scientific examination of science raises questions about the role of values in the scientific enterprise and how they might contribute to, as well as detract from, scientific decision-making. We close with a survey of contemporary application of probability and statistics to philosophical problems, followed by a sketch of some recent developments in the philosophy of physics, biology, and psychology.

In the last lecture, we finish bringing our two narratives together, and we bring some of our themes to bear on one another. We wrestle with the ways in which science simultaneously demands caution and requires boldness. We explore the tensions among the intellectual virtues internal to science, wonder at its apparent ability to balance these competing virtues, and ask how, if at all, it could do an even better job. And we think about how these lessons can be deployed in extra-scientific contexts. At the end of the day, this will turn out to have been a course in conceptual resource management.

Lecture One

Science and Philosophy

Scope: Standard first-lecture operating procedure would have me begin by trying to define *philosophy* and *science*, if not *of*. I think that's unwise at this point. Clarity and rigor, it is hoped, will be *results* of our inquiry, but we mustn't let them stand as forbidding *barriers* to inquiry. I try to dodge this problem, suggesting that relatively modest and uncontroversial characterizations of science and philosophy allow us to raise our central question, namely, what exactly is intellectually special about science. We then briefly examine some of the major epistemological and metaphysical issues raised by reflection on science. And we face, in a preliminary way, some important challenges to our enterprise. Does a scientific worldview leave any room for distinctively philosophical knowledge? And, more particularly, do philosophers really have anything useful to tell anyone, especially scientists, about science? Finally, we turn to the structure of the course, which involves a prequel, two long narratives, and a coda.

Outline

- I. Our classic way of beginning a lecture, especially a philosophy lecture, is by defining key terms. In this case, the key terms are *science* and *philosophy*.
 - A. But requiring a rigorous understanding of these notions right at the start makes it very hard to get going.
 - B. Major controversies arise about the nature of science and, even more so, about the nature of philosophy.
 - C. We will postpone detailed and controversial characterizations for as long as possible. All we need at the outset is a reasonably clear and simple statement of our central topic and some good reasons for getting interested in it.
- II. Our central topic is the special status of science. We'd like to understand why it's so special. And we can clarify this topic without resorting to elaborate or controversial definitions.
 - A. Science's most intriguing success is epistemic. We generally think that science is a good way to pursue knowledge—at least, about many questions. For this reason, it is natural to wonder what if anything unites the disciplines we call scientific and explains this distinctive epistemic success.
 - B. At the same time, our confidence in science is subject to significant limitations. There are many questions science cannot answer (at least for now) and many questions that it has answered incorrectly.
- III. But is philosophy the best place to try to discover what's epistemically special about science?
 - A. Many disciplines (such as history, sociology, and psychology) can make contributions to our understanding of what's distinctive about science.
 - B. Philosophy, in contrast, does not have its own domain of facts; thus, it's far from obvious what contribution philosophy can make to our understanding of science.
 - C. The best characterization I know of philosophy comes from one of my teachers: "Philosophy is the art of asking questions that come naturally to children, using methods that come naturally to lawyers."
 1. This leaves philosophy not only with its own fields, such as ethics, in which the childlike questions and lawyerly disputations have never gone out of style, but also with important intersections with scientific disciplines. Such questions as "What is space?" seem to belong both to philosophy and to physics.
 2. The question with which we began—what is so special about science?—is itself one of those bold, childlike questions that invites distinction-mongering and, thus, belongs more properly to philosophy than to any empirical discipline.
- IV. We can clarify this picture of the relationship between philosophy and science by contrasting it with two common and influential conceptions.
 - A. It was once widely believed that philosophy needed to serve as an intellectual foundation for the sciences.
 1. Real knowledge, it was thought, would have to be grounded in something more certain, more solid than observation and experience. Geometry served as a model, and almost all other disciplines fell short of that standard.

2. But philosophy's children have accomplished so much that they have changed the rules of the game and surpassed the intellectual prestige of their parent. Physics is now a paradigm of knowledge; philosophy is not.
- B. Does science, then, have any use for philosophy?
 1. All factual questions, one might think, are ultimately questions for some science or another. Any questions that are not scientifically answerable are, in some important sense, flawed.
 2. But this assertion sounds like a philosophical question, not a scientific one. The boundary between philosophy and other disciplines can be drawn only by doing philosophy. For this reason (among others), it's hard to avoid doing philosophy.
- V. A lot of good philosophy of science was done prior to the 20th century, but most philosophy of science these days is done in terms of a vocabulary and set of problems framed by the *logical positivists* (also known as *logical empiricists*; both terms emphasize the role of sensory experience in their views).
 - A. Though logical positivism is more or less dead, it figured centrally in the rise of philosophy of science as a unified subdiscipline. We will discuss the rise and fall of positivism through two main narratives.
 - B. We will begin, however, not with positivism but with the closely related views of the positivists' contemporary, Karl Popper. Popper offers the most influential approach to the most basic of our questions: What makes science *science*? His answer is very much *not* that scientific hypotheses are well supported by observational evidence.
 - C. We then approach positivism via Albert Einstein, the scientific hero of both Popper and the positivists. Einstein's work suggests that we have to be able to explain the meaning of our scientific terms by recourse to observation.
 - D. At this point, we'll be in a position to observe the positivists' struggle to develop the notion of the scientifically meaningful: Questions that go beyond experience in some ways are ipso facto unscientific (for example, whether humans have souls). But questions that go beyond experience in other ways (such as whether there are good reasons to believe in quarks) seem quintessentially scientific.
 - E. Along the way, we'll see that the positivists saw philosophy as akin to mathematics and logic and deeply different in methodology from the sciences. It aids the sciences by clarifying scientific concepts.
 - F. Staying within this broadly empiricist framework, we will turn from issues about observation and meaning to issues about observation and evidence. Can anything other than observational data count as evidence for the truth of a theory? How can there be a scientific method that allows us to go from relatively small observed samples to much grander conclusions about unobserved cases and unobservable objects?
- VI. Thomas Kuhn's work provides the first comprehensive alternative to the views of Popper and the positivists. Kuhn emphasizes the history of science, rather than its supposed logic.
 - A. Kuhn thought he could explain why science is a uniquely successful way of investigating the world without crediting science with being as rational, cumulative, or progressive as had been thought.
 - B. After presenting the essentials of Kuhn's work, we examine the reaction of two quite different groups of critics.
 1. One group held that, deprived of a special method, science can amount to only something like madness.
 2. The other group thought Kuhn insufficiently deflating of science's special epistemic status.
- VII. Having completed our first narrative, which primarily concerns meaning and evidence, we will return to positivism and take up scientific explanation and allied issues.
 - A. How can science explain while respecting its need to constrain itself within resources provided by experience?
 - B. Such notions as causation and physical laws likewise pressure science to go beyond the evidence of experience.
 - C. Finally, we ask about an especially ambitious and important kind of explanation: In what sense, if any, does the discovery of DNA allow genetics to "reduce to" molecular biology? And does biology itself reduce to physics?

- D. We will see how the tension between the ambitions of science to explain, to discover laws, and to unify disparate fields, on the one hand, and its insistence on confining itself within the bounds of experience, on the other, is resolved very differently by scientific realists than it had been by the logical positivists. This discussion will bring together aspects of our two major narratives.

VIII. The course closes with a two-part coda. We examine the probabilistic revolution that has made such a difference to the recent philosophy of science, asking how that allows us to reframe issues of objectivity and justification. And we end by looking at examples from within philosophy of physics, biology, and psychology to apply what we have learned in the general philosophy of science and to examine some of the philosophical issues that arise within particular sciences.

Essential Reading:

Rosenberg, *Philosophy of Science: A Contemporary Introduction*, chapter 1.

Godfrey-Smith, *Theory and Reality: An Introduction to the Philosophy of Science*, chapter 1.

Supplementary Reading:

Hitchcock, *Contemporary Debates in Philosophy of Science*, introduction.

Questions to Consider:

1. This lecture suggests that the claim that science can settle all factual questions is a philosophical, not a scientific, thesis. Why is that? What makes a thesis philosophical?
2. What shifts in intellectual values had to take place for science to surpass philosophy in cultural prestige?

Lecture Two

Popper and the Problem of Demarcation

Scope: Now we can get serious about what science is. Can we distinguish, in a principled way, between sciences and pseudosciences? We often talk as if even quite unsuccessful scientific theories deserve a kind of respect or standing that should not be accorded to pseudoscientific theories. Inspired by Einstein's work, Karl Popper offers a striking, elegant, and influential criterion for distinguishing genuine from counterfeit science. Popper denies the seemingly obvious claim that scientists seek highly confirmed theories. The distinguishing mark of science, for Popper, is that it seeks to falsify, not to confirm, its hypotheses. In this lecture, we develop and assess this remarkable proposal. Can Popper sustain the claims that his examples of pseudosciences fail his test and that his examples of genuine sciences pass it? Could science function effectively if it were as open-minded as Popper says it should be?

Outline

- I. The problem of demarcation challenges us to distinguish, in a motivated and non-arbitrary way, between genuine sciences and pseudosciences.
 - A. Not every non-science is a pseudoscience. A pseudoscience is a discipline that claims the special epistemic status that science holds for the same reasons that science makes that claim but does not, in fact, merit that status.
 - B. To call something a pseudoscience is not to deny that it might sometimes make true and important claims. Likewise, to call something scientific is not to deny that it might well be false. Scientific claims, we tend to think, merit a kind of consideration to which pseudoscientific claims are not entitled.
 - C. The problem of demarcation is of clear practical, as well as theoretical, importance.
 - D. It would be nice to have a clear definition of *science*, but a good deal of progress can be made without reaching a definition.
- II. Karl Popper's elegant solution to the demarcation problem has been enormously influential, especially among scientists.
 - A. Popper's theory arises from the intellectual context in which he (along with the logical positivists) came of age.
 1. Popper was especially interested in Einstein's theory of relativity, Karl Marx's theory of history, and the psychological theories of Sigmund Freud and Alfred Adler.
 2. It was widely believed at the time that the work of Marx, Freud, and Adler was genuinely scientific, but Popper became disenchanted with such theories.
 3. Popper argued that Einstein's theory was distinguished from those of Marx, Freud, and Adler by its openness to criticism. This provides the key to Popper's solution to the problem of demarcation.
 - B. Popper's emphasis on criticism stems from his rejection of the most straightforward criterion of demarcation, according to which scientific claims are special because they are confirmed by observational evidence and because they explain observations.
 1. Pseudosciences, such as astrology, are chock full of appeals to observational evidence. Observation, for Popper, is cheap. It is essentially interpretation of experience in terms of one's theory. The pseudoscientist finds confirming evidence everywhere (for example, in the many case studies of Freud and Adler).
 2. Furthermore, apparent counterevidence can be turned aside or even turned into confirming evidence by a clever pseudoscientist. Freud and Adler had ready explanations for any observational result.
 3. For Popper, no evidence falsifies a pseudoscientific claim and almost everything confirms it. As a result, Popper came to see the two standard virtues of scientific theories—explanatory power and confirmation by a large number of instances—as closer to being vices than virtues.
 4. Fitting the data well is, thus, not the mark of a scientific theory; a good scientific theory should be informative, surprising, and in a certain sense, improbable.
 - C. Einstein's theory of relativity, on the other hand, came to exemplify genuine science for Popper.

1. General relativity led to the surprising prediction that light would be bent by the gravitational field of the Sun. It was a great triumph when Arthur Eddington's expeditions verified that light was bent by the amount that Einstein had predicted.
2. For most observers, what mattered was the fit between Einstein's predictions and the evidence, but not for Popper. What mattered to him was that the theory had survived a severe test. The mark of a genuinely scientific theory is falsifiability. Science should make bold conjectures and should try to falsify these conjectures.

III. Popper's theory is admirably straightforward, but it nevertheless requires some clarification.

- A. Popper generally writes as if falsifiability and, hence, scientific standing come in degrees. This suggests, however, that pseudosciences differ more in degree than in kind from genuine sciences.
- B. Popper's theory is both descriptive and normative. He claims both that this is what scientists do and that it is what they should do.
- C. Popper is not offering a definition but only a necessary condition. He is not saying that all falsifiable statements are scientific but only that all scientific statements are falsifiable. Falsifiability is a pretty weak condition.
- D. To call something unscientific is not to call it scientifically worthless.
 1. Popper thought that Freud, Marx, and Adler said some true and important things.
 2. Furthermore, metaphysical frameworks, such as atomism (which was not testable for centuries after it was proposed), can help scientists formulate testable hypotheses.
 3. Popper even thought for awhile that Darwin's principle of natural selection was an ultimately unscientific doctrine. He later changed his mind about this, arguing that the Darwinian claim about survival of the fittest is not a mere definition of fitness (and, hence, unfalsifiable) but instead implies historical hypotheses about the causes of traits in current populations.

IV. Popper's view faced some serious criticisms.

- A. Such statements as "There is at least one gold sphere at least one mile in diameter in the universe" do not seem to be falsifiable on the basis of any finite number of observations, but they do not seem unscientific either. More important, statements involving probabilities appear unfalsifiable. A run of 50 sixes in a row does not falsify the claim that this is a fair die.
- B. Popper does not adequately distinguish the question of whether a theory is scientific from the question of whether a theory is handled scientifically. Are theories scientific in themselves or only as a function of how they are treated?
- C. Good scientific theories aren't cheap. It is not clear that scientists do or should reject theories whenever they conflict with observed results.
- D. Should we accept the idea that being highly confirmed and having wide explanatory scope are not virtues of a scientific theory? Was it not a striking feature of Newton's physics that it could explain the tides, planetary motion, and so on?
- E. Thus, it is not exactly clear how Popper's view should be expressed: Is it about the logical form of scientific statements or about the way they are treated by their advocates? However it is formulated, it is not clear that it provides a necessary condition for science.

Essential Reading:

Popper, "Science: Conjectures and Refutations," in Curd and Cover, *Philosophy of Science: The Central Issues*, pp. 3–10.

Supplementary Reading:

Kuhn, "Logic of Discovery or Psychology of Research?" in Curd and Cover, *Philosophy of Science: The Central Issues*, pp. 11–19.

Questions to Consider:

1. Is there a better way to characterize observation than “interpretation in the light of theory”?
2. Can you describe conditions under which you think scientists would reject central and widely accepted hypotheses (such as the fundamentals of evolution by natural selection or of plate tectonics)? How significant is the ease or difficulty with which you accomplish this task?

Lecture Three

Further Thoughts on Demarcation

Scope: Given the enormous practical importance of demarcating science from pseudoscience, it comes as no surprise that Popper's criterion has competitors as well as critics. We survey a number of proposals and see how they apply to (allegedly) clear cases of science, (allegedly) clear cases of pseudoscience, and more controversial cases, such as creationism. Though many contain valuable insights, no demarcation criterion has won widespread assent, and we take stock of this situation. What would be the implications of deciding that astrology is better described as lousy science than as pseudoscience? Would this inevitably lead to the teaching of creationism in high school classrooms?

Outline

- I. The issue of falsifiability (or, more generally, testability) is a tricky one, and its slipperiness is one of the major reasons philosophers have not generally found Popper's approach to demarcation persuasive. It is difficult to interpret Popper's falsificationism so that physics passes the test and Freud, for example, fails it.
 - A. Often, a pseudoscientist makes predictions that are admitted to be false, but the theory is not taken to be falsified. It is crucial to realize that a false prediction is not a sufficient basis for rejecting a theory. Complex sciences, such as medicine, tolerate quite a number of false predictions.
 - B. We cannot require that a theory be rejected (either as bad science or as pseudoscience) merely because of persistent failures of fit with the evidence. We would have little science left; much scientific work involves trying to resolve these failures of fit.
 - C. But neither can we simultaneously reject a theory for making false predictions and for failing to make falsifiable predictions.
 - D. My claim is not that there's no difference between astrology and physics with respect to falsifiability, but only that this difference is surprisingly hard to characterize.
- II. What other demarcation criteria do we have? One interesting criterion is historical: Pseudosciences tend not to make much progress.
 - A. But progress can be tricky to characterize, much less to measure.
 1. Astrology has certainly changed over the centuries, and it's plausible to claim that some of the changes constitute improvements.
 2. A science that correctly accounted for everything in its domain could hardly be expected to show much progress.
 - B. A more sophisticated version of this approach might fault a pseudoscience in comparison to rival theories. If a competitor makes substantial progress while the theory in question remains stagnant, then the unprogressive theory becomes pseudoscientific.
 1. This view has the consequence that a theory's scientific status can change over time, without any change in the theory itself.
 2. More troublingly, this criterion appears to have the consequence that theories that lack serious competitors are not pseudosciences.
- III. Several other criteria have been put forward, but each of them seems, at best, problematic.
 - A. Pseudosciences, such as astrology, often lack a clear mechanism; no explanation is offered of *how* the stars influence our lives. But many legitimate and successful theories lack mechanical accounts of crucial processes. Isaac Newton provided no physical mechanism for the action at a distance of gravity, for instance.
 - B. Some adopt a kind of social practice conception of science. A practice counts as scientific if the right people call it a science (and if its practitioners do the right sort of scientific things, such as publish journals and get jobs in universities). But this criterion counts institutionalized pseudoscience (for example, Lysenkoist biology) as scientific.

- C. Many pseudosciences have epistemically dubious origins, but genuine sciences, including chemistry, also originated in such dubious enterprises as alchemy, and almost all science ultimately arose from mythology and speculation.
 - D. Nor do there seem to be forms of reasoning that distinguish science from pseudoscience.
 - 1. Pseudosciences appear to use mathematical reasoning and to make causal and explanatory inferences.
 - 2. Genuine sciences sometimes use more hazardous forms of reasoning, such as arguments from analogy and other strategies that figure prominently in pseudosciences.
- IV. Creationism occasions the most heated debates about demarcation.
- A. Young-Earth creationism (YEC) makes relatively specific assertions about the creation of the universe from nothing, the age of the Earth, and about the separate creations of “kinds” of creatures.
 - B. Intelligent-design creationism (IDC) refrains from making claims as specific as those put forward by YEC. Intelligent-design theorists focus on what they consider the core creationist principles, to wit, that there is a personal, supernatural creator of the universe who continues to influence creation and does so for some purpose.
 - C. YEC and IDC can unite on certain negative arguments against Darwinism and, perhaps, against other parts of the “naturalistic worldview.” What is the scientific status of these arguments?
 - 1. The negative arguments concern such matters as the limitations of the fossil evidence for evolution and the supposed inability of natural processes to account for certain kinds of complexity.
 - 2. Can such negative arguments suffice for scientific status? On the one hand, it seems plausible that one could spend a valuable scientific career doing nothing but research aimed at falsifying, say, the wave theory of light. On the other hand, there is surely no scientific discipline called “the wave theory of light is wrong.”
 - 3. Thus, if we’re asking about YEC and IDC as disciplines, it is plausible to insist that their status depends, at least in part, on the status of their positive proposals. Demarcation might apply differently to the work of individuals, however.
- V. YEC has not fared well in the American court system; it has generally been pronounced pseudoscientific there. What are the arguments for this conclusion and how good are they?
- A. One common complaint is that this theory explicitly invokes supernatural causes and, thereby, disqualifies itself as scientific. This complaint won’t get much traction unless the natural/supernatural distinction can be drawn independently of the scientific/unscientific distinction.
 - B. It might be true and important that YECists refuse to treat any evidence as falsifying their theory. But we must distinguish criticisms of the proponents of theories from criticisms of the theories themselves. Would a group of physicists’ refusal to treat any evidence as falsifying quantum mechanics show *the theory* to be unscientific?
 - C. Similarly, most YECists would admit to having religious motivations for their work. But many scientists have been motivated by religious beliefs, and some scientists are motivated by money. In none of these cases do the motives render the work unscientific.
 - D. YEC explanations make relatively little use of natural laws and mechanisms. But some scientific theories make little use of laws and/or lack crucial mechanisms.
 - E. From the standpoint of mainstream science, anyway, claims by YEC about the age of the Earth are testable (and false).
- VI. IDC theorists have offered a much thinner research agenda than YEC proponents have, and this raises quite different demarcation questions.
- A. IDCists argue, quite plausibly, that there need be nothing unscientific about the search for intelligent design. Many scientists have thought it plausible that we could get evidence of extraterrestrial intelligence.
 - B. The next step in the main IDC argument is the crucial one. It claims that certain kinds of complexity found, for instance, in earthly organisms are thought to provide evidence of intelligent design. This is very like the classic “design argument” for God’s existence.

- C. We're asking whether the argument is scientific, not whether it is strong. One major problem is that IDC seems dominated by big questions, and it doesn't seem to have much going on in the way of little questions that can be answered in labs.

VII. Most philosophers think that the demarcation problem has not received an adequate solution.

- A. The notion of demarcation might not apply univocally to theories, to individuals, and to disciplines.
- B. We haven't seen a solid basis for distinguishing between poor scientific theories and nonscientific theories.
- C. If the classic demarcation project is abandoned, it won't be possible to say that creationism (or astrology) is unscientific. But if that's the case, qualifying as scientific won't be much of an accomplishment.
 - 1. Should we decide which theories receive funding and which are taught in schools on the basis of which theories are good, rather than which theories are scientific? Of course, we'll need criteria of goodness (see the rest of the course).
 - 2. The legal and political issues raised here (for example, the Constitution does not forbid teaching bad science, assuming for the sake of argument that creationism constitutes bad science) are beyond the scope of our course.
- D. From the fact that no adequate demarcation criteria have been formulated, it doesn't follow that none can be formulated.

Essential Reading:

Thagard, "Why Astrology Is a Pseudoscience," in Curd and Cover, *Philosophy of Science: The Central Issues*, pp. 27–37.

Exchange between Ruse and Laudan on creation science in Curd and Cover, *Philosophy of Science: The Central Issues*, pp. 38–61.

Supplementary Reading:

Pennock, ed., *Intelligent Design Creationism and Its Critics: Philosophical, Theological and Scientific Perspectives*.

Questions to Consider:

- 1. Justice Potter Stewart famously said that though he couldn't define pornography, he knew it when he saw it. To what extent are you confident that you know pseudoscience when you see it?
- 2. How do you think that the legal issues surrounding evolution and creationism would change if we gave up trying to find a demarcation criterion? The U.S. Constitution (arguably) forbids the teaching of religion, but it doesn't seem to ban the teaching of less-than-stellar science. Even if Darwinists could show that evolutionary biology is (at least for now) a better theory than intelligent design, could the latter view legitimately be banned from public school classrooms?

Lecture Four

Einstein, Measurement, and Meaning

Scope: Einstein's special theory of relativity delivered a shock to physicists and to scientifically minded philosophers. Relativity didn't just point out surprising new facts, and it didn't merely require strange new concepts. It revealed a disturbing lack of clarity lurking within familiar concepts, such as those of length and simultaneity. Einstein's work suggested that physics (and philosophy) had been working with an inadequate conception of concepts. Though he did not offer it as a demarcation criterion, the philosophically inclined Nobel laureate P. W. Bridgman proposed an influential theory by which scientific concepts must be expressed in strongly experiential terms. Bridgman's *operationalism* faced serious problems, but it leads us nicely into a discussion of science as distinguished from other enterprises by the way in which it *disciplines* its conceptual and evidential resources in the light of experience.

Outline

- I. In order to understand why Einstein's special theory of relativity exerted such influence on philosophers of science, we need to understand the central problem that Einstein solved.
 - A. We are reasonably familiar with the idea that unaccelerated motion can be detected and described only with respect to some reference frame. This leads to something worth calling a *principle of relativity* (though it long predates Einstein). If two people float past each other in the depths of empty space, there is no way to tell which of them is really moving. It is tempting to say that the question of which one is really moving has no meaning.
 - B. On the other hand, there was some reason to think that sense could be made of something rather like absolute motion by reflecting on light.
 1. James Clerk Maxwell (writing in the mid-19th century) had shown light to be a kind of electromagnetic wave. It was generally believed that light moved through a pervasive aether. And a reference frame at rest with respect to the aether (which pervaded space) would be pretty close to the reference frame of space itself.
 2. If the world were as 19th-century physics took it to be, we would be able to measure our motion through the aether by detecting differences in the observed speed of light. We would be catching up to the light in one direction (so it should appear to move more slowly than it would to an observer at rest in the aether) and running away from it in another (in which case, the opposite would happen).
 3. But experiments failed to detect any motion of the Earth with respect to the aether. Experiments consistently measured the same speed for light in all directions (just as would be expected if one were always at rest with respect to the aether). Light seemingly disobeyed the "all (unaccelerated) motion is relative" slogan.
 - C. The two principles associated with Einstein—the relativity of all (unaccelerated) motion and the stubborn unrelativity of the speed of light—seemed to contradict each other.
- II. Einstein overcame the apparent tension between these principles by critically examining some of our most central concepts. The principles contradict each other only if certain assumptions about space and time are in place.
 - A. When combined, the principles imply that observers moving relative to one another will, if all their instruments are sufficiently sensitive and functioning properly, get different answers to such questions as whether one event happened before another.
 - B. These seemingly incompatible observations can all be correct only if there is something wrong with such questions as "When did event E happen?" Einstein suggests that such questions are scientifically meaningless unless a reference frame is specified.
 - C. Similar considerations apply to the measurement of space. Observers in motion with respect to one another will measure the length of an object differently. All can be right, provided we reject the notion that the object's length is independent of the reference frame from which it is measured.

- D. Other physicists were unable to reconcile the experimentally established principles because they assumed that they had a clear understanding of such concepts as *simultaneity* and *length*. Much of Einstein's achievement involved linking such concepts very tightly to experience and measurement, while denying that they had legitimate use when disconnected from experience and measurement. This idea exerted enormous influence on physicists and philosophers.
- III. We can now turn to more directly philosophical matters and begin exploring a question that will occupy us for some time: In what way must a concept be "cashed out" in experiential terms in order to count as scientifically legitimate? P. W. Bridgman provides the most directly Einstein-inspired example.
- A. Never again, says Bridgman, are concepts to prevent us from seeing what nature tries to show us. The way to prevent this is to be sure that something in nature answers to each of our concepts. And the way to do that, according to Bridgman's *operationalism*, is to define each scientific concept solely in terms of the operations required to detect or measure instances of the concept. Thus, length is to be identified, not with some property, such as taking up space, but with the procedures for using a meter stick. This is all that *length* means.
- B. Strictly speaking, each operational procedure generates a distinct concept, for example, *alcohol-thermometer temperature* and *mercury-thermometer temperature*. Officially, we change the subject whenever we change procedures, because the procedure is the meaning. Bridgman wants to make us aware of the risk we run when we assume that these two concepts refer to the same physical magnitude.
- C. We need a basic vocabulary in which operational definitions can be given. Operations have to end at something that does not require further operationalizing. Bridgman assumes that some phenomena are directly and unproblematically observable and, thus, not in need of operational definition.
- IV. Operationalism has been enormously influential in many scientific disciplines, but many philosophers think operationalism represents a too-stringent way of tying down our concepts in experiential terms.
- A. Operationalizing weight in terms of a pan balance assumes that no "additional" forces are affecting the pans differently. But how are we to specify "no additional forces" in observational and/or operational terms?
- B. Our confidence that two different kinds of thermometers measure the same "stuff" relies on an idea of the thing being measured that far outruns the measurements. If we were trying to build a device that would measure the temperature of the Sun, we'd be relying on the notion of a good temperature-measuring device. But at that point, we have given up reducing the notion of temperature to what we can actually measure, and that was supposed to be the point of Einstein's story.

Essential Reading:

Greene, *The Elegant Universe: Superstrings, Hidden Dimensions, and the Quest for the Ultimate Theory*, chapter 2.
 P. W. Bridgman, "The Operational Character of Scientific Concepts," in Boyd, Gasper, and Trout, *The Philosophy of Science*, pp. 57–69.

Supplementary Reading:

Sklar, *Philosophy of Physics*, chapter 2.
 Hempel, "A Logical Appraisal of Operationism," in Brody and Grandy, *Readings in the Philosophy of Science*, pp. 12–20.

Questions to Consider:

1. Many philosophers and physicists felt Einstein's revolution to be a distinctively conceptual one. Does this seem right to you? Newton's and Darwin's revolutions certainly involved far-reaching conceptual changes. Why, if at all, does special relativity count as an especially conceptual scientific shift?
2. How can an operationalist make sense of the idea that a measuring device (such as a thermometer) is malfunctioning?

Lecture Five

Classical Empiricism

Scope: In order to develop a more sophisticated understanding of the connections between experience and meaning than operationalism can provide, we need to draw on a rich history of philosophical reflection about experience, language, and belief. John Locke, George Berkeley, and David Hume constitute a tradition united by its *empiricism*—the idea that experience sets the boundaries of, and provides the justification for, our claims to knowledge. We will examine classic empiricist analyses of matter and mind and see that empiricism’s admirable anti-metaphysical tendencies constantly threaten to force it into a disabling and radical skepticism. In fact, we will see that classical empiricism has difficulty making room for the possibility of classical empiricist philosophy. The classical tradition sets the terms of the problems that a sophisticated empiricist account of scientific knowledge will have to solve.

Outline

- I. Einstein and Bridgman were philosophically inclined physicists. The problem with which they were wrestling, that of how concepts have to be connected to experience to be legitimate, has a long philosophical history. Systematic philosophical reflection about experience as a source of and constraint on our knowledge really begins with John Locke.
 - A. For this reason, Locke is often considered the first empiricist (*empiricism* is roughly the view that sensory experience is the ultimate source of our concepts and of our knowledge).
 - B. Locke’s project most directly concerns knowledge: He wanted to determine the boundaries of human knowledge.
 - C. Locke investigated the scope of our knowledge by investigating its sources. He claimed that experience is the source of all the material of thought: “Nothing is in the mind that was not first in the senses.”
 1. An idea, for Locke, is what is in the mind when the mind thinks. Ideas are mind-dependent; they are (more or less) literally in minds. The things I directly perceive are sights and sounds, not physical objects.
 2. Simple ideas are given in experience. Innate mental powers (notably combination and abstraction) allow us to refine and extend our simple ideas. Abstraction lets us focus on a part of a presented idea (for example, the blueness of the sky), and these parts can be recombined to form ideas of things never presented in experience, such as unicorns.
 - D. Locke recognized the limitations of what experience puts us in a position to know. We have very little understanding of the inner nature of material substances, and we are unable to form any useful idea of how such substances produce in us many of the ideas they generate.
 - E. Locke’s highly influential view represents something of a standard empiricist bargain. We gain systematic resources for clarifying our ideas, and we pay for this clarification by realizing that we don’t get to know as much or even say as much as we might have thought we could.
- II. Though an empiricist himself, George Berkeley’s work suggested that the conceptual costs we pay for confining ourselves to what is presented in experience are much more radical than Locke thought.
 - A. Berkeley saw himself as purging philosophy of its tendencies toward skepticism and atheism, but he was much misunderstood by his contemporaries.
 - B. It is perhaps understandable that his contemporaries thought him a skeptic, because Berkeley denied the existence of matter. A material object is supposed to be something that “holds” or “supports” its properties, and Berkeley goes so far as to deny that we have an idea of material substance.
 1. We have no direct experience of matter. What does it look or feel like?
 2. Berkeley denied that we can obtain a legitimate idea of matter through abstraction. We cannot imagine a thing without its properties.
 3. Locke had already admitted that it was mysterious how material objects produced ideas in us.
 - C. For Berkeley, God simply produces ideas in us directly. God does not use matter as an intermediary way to cause our experiences.

- D. As a result, Berkeley was the first empiricist to get over the idea that we need to get behind or beyond experience.
 - 1. For Berkeley, the patterns in our experience are the world itself. God has set things up so that if we formulate and apply, say, Newton's laws of motion, we can predict what experiences we will have.
 - 2. All science can or should be is the development of rules for predicting what experiences we will have.
- III. It took an empiricist of the next generation, David Hume, to show how devastating the skeptical consequences of a resolutely pursued empiricism can be.
 - A. Hume's project is not itself skeptical. He aspired to bring the "experimental method" to bear on philosophy.
 - B. But a rigorously applied experimental method finds that many crucial notions do not have a proper pedigree in experience (in Hume's lingo, we have no impressions answering to such notions).
 - 1. Hume held that we have no impression of causation, of one event making another event happen. All experience shows us is one thing after another. The connections between them are not experienced.
 - 2. We have no impressions of enduring things. Our experience is constantly changing; the sensations we have do not endure and are not constant.
 - 3. Nor do we have impressions of ourselves as things that endure through time. We are not thinking things but bundles of impressions.
 - 4. Experience provides us with no clear concept and nothing worth calling evidence for the existence of anything not currently perceived by us. This is very deep skepticism indeed.
 - C. As a result, many of our most basic notions are either meaningless or have very different meanings than we might have thought they had. My idea of myself, for instance, is cobbled together by the imagination, rather than by reason or experience. We are much less reasonable than we think we are (and it's a good thing, too!).
 - D. Hume faces a philosophical problem about philosophy more squarely than his predecessors had. Where can philosophy fit into an empiricist framework?
 - 1. Hume held that all meaningful statements must concern either relations of ideas, as in logic and mathematics, or matters of fact, as in the empirical sciences. This influential dichotomy is known as *Hume's fork*.
 - 2. Hume saw himself as addressing matters of fact. He thought that he was doing a kind of psychology, seeking the laws that govern the mind, as Newton had sought and found the laws governing nature.
 - 3. But is Hume really doing psychology? If philosophy is not psychology without the experiments, what might it be?
- IV. We leave the 17th and 18th centuries with two challenges for later empiricists.
 - A. Is there a way to reconcile the core empiricist idea that experience is the source of our conceptual and evidential resources with the apparent need to go beyond what is presented in experience if we are to do science or philosophy?
 - B. Does philosophy connect to experience in the right sort of way to be a legitimate discipline? Is philosophy just science with low evidential standards?

Essential Reading:

Berkeley, *Three Dialogues between Hylas and Philonous*.
 Hume, *A Treatise of Human Nature*, Book 1.

Supplementary Reading:

Woolhouse, *The Empiricists*, especially chapters 6–8.

Questions to Consider:

1. Locke argues that we lack the sensory capacities that would be required to know the real nature of such substances as gold. But many people think that, despite our limited sensory capabilities, we have attained knowledge of the real nature of gold. Has Locke's argument gone wrong, and if so, where?
2. Most of us think, *pace* Berkeley, that we do have a legitimate idea of matter. If so, from where does it come? Is it innate? If it arises through experience, how does it do so?

Lecture Six

Logical Positivism and Verifiability

Scope: Like Popper's philosophy of science, *logical positivism* (also known as *logical empiricism*) was born in the first decades of the 20th century in the German-speaking world. Like Popper, the positivists were inspired by Einstein's stunning successes. But unlike Popper, they were deeply interested in classic empiricist questions about the connections between meaning and experience. Drawing on recent developments in logic and the philosophy of language, they tried to develop an empiricist conception of philosophy that was logically coherent and adequate to the practice of science. In this lecture, we motivate and sketch the positivist program, paying special attention to their demarcation criterion, the (in)famous *verification principle*.

Outline

- I. The logical positivists made philosophy of science a major subfield for the first time. Their approach to the field dominated for decades.
 - A. They were highly impressed by Einstein's work and other developments in physics and highly unimpressed by much of 19th- and early-20th-century German philosophy. To them, the philosophy of the day seemed like armchair speculation, much of which stood in the way of scientific progress.
 - B. They were less worried than Popper was about pseudosciences and more worried than he was about metaphysics and about philosophy getting in the way of physics. This leads, as we will see, to a different approach to the demarcation problem.
 - C. The *positivism* part of logical positivism derived from the 19th-century French thinker Auguste Comte and reflects his animus against traditional metaphysics.
 - D. The *logical* part of logical positivism reflects the positivists' belief that mathematical logic provided tools with which a new and improved version of empiricism could be built, one that would be favorable to science and unfavorable to metaphysics.
 - E. This new version of empiricism grasped the other option presented by Hume's fork. For the positivists, the philosopher deals in relations of ideas, not matters of fact. Philosophy clarifies linguistic problems and exhibits the relationships between scientific statements and experience.
- II. The basic principle of the positivist program states that every *cognitively meaningful* statement is either *analytic* or is a claim about possible experience.
 - A. Cognitively meaningful statements are those that are literally true or false.
 1. Imperatives and questions have meaning but are not statements in the relevant sense. They are not candidates for truth.
 2. We find statements in poetry, but they, likewise, do not aim for literal truth.
 - B. Analytic statements concern Hume's relation of ideas. They are true or false in virtue of their meanings and have no factual content.
 1. Consequently, they are *knowable a priori*; we do not need empirical evidence in order to know the truth of logical and mathematical propositions.
 2. Analytic truths also hold *necessarily*. It is not merely true that no bachelor is married; it must be true. Such a statement is "true in all possible worlds."
 3. We can, thus, be certain that every effect has a cause, but this is no great metaphysical insight; it is, rather, a fact about how we use the words *cause* and *effect*.
 - C. A traditional metaphysical statement is one that has factual content (that is, one that is synthetic, not analytic) yet is supposed to be knowable independently of experience.
 1. Such statements purport to make factual claims that are supposed to hold no matter what experience seems to show (for example, "Every event has a cause"). But the content of a factual statement, according to positivism, is exhausted by what the statement says about possible experience.

2. Thus, metaphysical statements are not false; they are not candidates for being true or false. At best, they are unintentional poetry.

III. How are we to tell when we are dealing with meaningful statements?

- A. The logical positivists talked of meaningfulness in terms of verification. To be cognitively meaningful is to be either true or false; thus, a statement is meaningful if there is the right sort of method for testing truth or falsehood.
- B. For analytic statements, the model is mathematical or logical proof. Thus, analytic statements are verifiable and, hence, meaningful if they can be traced back in the appropriate way to their source in linguistic convention.
- C. Our main concern is with empirical (that is, synthetic) statements.
 1. Where operationalism and classical empiricism focus on the connections of a *term* to experience, the verificationism of the positivists makes empirical meaningfulness a matter of a *statement's* ability to confront experience.
 2. This represents a significant liberalization of empiricism (made possible, in part, by advances in logic). A term can get its meaning from its role in making meaningful statements; it need not be established as independently meaningful.
- D. The verifiability of a synthetic statement involves finding *possible* observations that bear on its truth.
 1. If we required *actual* observations, we would be assessing the truth or falsehood of the statement, not its meaningfulness.
 2. The sense in which such observations must be possible presents difficult problems.
- E. The most straightforward way to be sure that a statement is verifiable would be to determine a set of possible observations that would conclusively show the statement to be true.
 1. But this is too demanding. No finite number of observations could conclusively establish the truth of "All copper conducts electricity."
 2. Similar problems face a broadly Popperian proposal that substitutes conclusive falsifiability for conclusive verifiability. Even combinations of these two proposals face counterexamples.
 3. Perhaps most importantly, statements about unobservable objects appear to get ruled out by this criterion. How could observations conclusively establish that "That streak in the cloud chamber was produced by an electron"?
- F. For this reason, we need a weaker version of the verifiability principle.
 1. A. J. Ayer suggested that if we can use the statement to derive observation statements that cannot be derived without it, the statement is meaningful.
 2. But this is much too weak, because it does not impose any restrictions on the *auxiliary hypotheses* we can use in our derivation. From "Everything proceeds according to God's plan" and "If everything proceeds according to God's plan, then this litmus paper will turn pink when placed in this solution," it is easy to derive an observational prediction. We need the statement about God's plan to do the derivation.
 3. Ayer modified his principle to try to require that the auxiliary hypotheses be independently meaningful, but this proposal succumbs to technical objections.
- G. Perhaps surprisingly, positivism was not derailed by the difficulties involved in formulating an adequate version of the verifiability principle. The idea that empirical meaningfulness had to get construed in terms of observation remained powerful, though it resisted clear encapsulation.

Essential Reading:

Ayer, *Language, Truth and Logic*, especially the introduction and chapters I–III.

Supplementary Reading:

Godfrey-Smith, *Theory and Reality: An Introduction to the Philosophy of Science*, chapter 2.

Soames, *Philosophical Analysis in the Twentieth Century*, chapters 12–13.

Questions to Consider:

1. Do metaphysical statements, such as “Every event has a cause” and “Human beings have free will,” seem (cognitively) meaningless to you? Can you account for such meaning as you think such statements have within the framework of positivism?
2. If very few statements can be conclusively verified or conclusively falsified, then few statements can be *proved* on the basis of experience. But we often talk of experimental proof. Is such talk exaggerated?

Lecture Seven

Logical Positivism, Science, and Meaning

Scope: Having looked in a general way at the positivist requirements for meaningfulness, we now turn our attention directly to scientific theories. As we have seen, empiricism has trouble with unobservables—it is difficult for an empiricist to make room for intelligible talk, much less knowledge, of unobservable reality. But scientific theories are chock full of claims about quarks and other apparently unobservable entities, and they also invoke dispositions (like solubility) and other suspiciously metaphysical-sounding properties. Attempts to reduce talk of unobservables to talk of observable reality appear to be too stringent, while more permissive attempts to reconcile the demands of empiricism with the importance of unobservables in science threaten to allow metaphysical statements to count as meaningful. A key consequence of all this empiricism is *instrumentalism*, according to which a scientific theory need only “save the phenomena.”

Outline

- I. Logical positivists needed to show, as their empiricist predecessors had not, that science could be adequately reconstructed in empiricist terms.
 - A. The logical positivist conception of how scientific theories work was so influential that it is generally called the “received view of theories.”
 - B. Unsurprisingly, given the logical positivists’ conception of the business of philosophy, they thought of a scientific theory as a linguistic kind of thing. It is a set of sentences that has certain properties.
 1. For purposes of explicitness and clarity, they envisioned theories stated in the language of logic.
 2. They were not saying that this is the best form for doing science; rather, it is the best form for displaying the relationships of meaning and evidence that make science special.
 3. This is a distinctive approach to science called a *rational reconstruction*.
 - C. The language of logic presents no problems of meaningfulness. But you need more than just logical connectives in order to do science. We need to be able to give an empirical interpretation of such language as “There is an object X such that X has property P.”
 1. We can help ourselves to terms that refer to observable objects and properties. The positivists, like their classical empiricist predecessors, take such terms to be unproblematically meaningful.
 2. But we’re not going to be able to do any science on the basis of observational and logical vocabularies alone. We can list observations, but we will not be able to do any predicting or explaining, and that is the heart of science.
 3. Our theories need theoretical terms, such as *acid* and *litmus paper*, if we are going to have any scientific understanding of the world. But none of these terms belongs in the logical or the observational vocabulary.
 4. This encapsulates a huge, recurring tension: Science must limit itself to experience and it must go beyond experience.
 - D. How are we to expand the vocabulary without violating empiricism and opening the door to metaphysics? We can try to explicitly define new terms on the basis of already legitimate terms.
 1. We would like to use, for example, *fragile* to predict and explain things. But this is not an observation term. You cannot tell just by looking whether something is fragile. You have to whack it. *Fragile* is a *disposition term*; it refers to a property that manifests itself only under certain test conditions.
 2. We cannot define “X is fragile” as “If we strike X, it will break.” That has the consequence that anything we fail to strike is fragile.
 3. We want to define “X is fragile” as “If we were to strike X, it would break.” But this *counterfactual conditional* cannot be defined in terms of the logical vocabulary or the observational vocabulary. Such conditionals depend on messy facts about how the world would be if were different than it actually is.
 - E. We can retreat to *partial definitions* of new terms in the observational and logical vocabulary.
 1. What we can say is something like this: “Anything struck [with a ‘standard’ whack] is fragile just in case it breaks.” This statement is only a definition of fragility for struck objects; it refuses to commit itself to anything about the fragility of unstruck objects. For this reason, it is a partial definition.

2. A partial definition has empirical content. We can use partially interpreted terms to make predictions (for example, that a piece of crystal will break when struck).
- F. But we have to keep moving away from the observational level in order to explain phenomena and in order to generate predictions about more complex phenomena. For instance, *fragility* will need to be hooked up to claims about molecular structure or something similar if it is going to be of any real scientific interest.
1. Thus, we need to keep expanding the vocabulary, partially interpreting new terms on the basis of other partially interpreted terms. We need statements linking terms in the new “theoretical” vocabulary “down” to observation and “up” to statements and terms that stand at an even greater remove from observation.
 2. A scientific theory is structured like a mathematical theory, with the most general laws serving as axioms. The most fundamental laws, such as Newton’s laws of motion, provide the theory’s basic explanatory framework.
 3. Empirical meaning comes in via those statements of the theory that directly connect to observation, and the deductive relationships among the theory’s statements serve to spread that meaning around the theory. This is the received view of theories.
- II. But once we think about how complex and far removed from experience many scientific claims are, there’s a danger that we’ve lost track of anything worth calling experiential meaning. By loosening up the strictures to allow for realistic science, there’s a danger that we will have let in metaphysics.
- A. What stops me from introducing the following partial definition into my chemical theory: “A sample of water is ‘unholy’ if it has ever been used to make light beer”? This allows me to predict some places where unholy water will be found.
 - B. The classic response is that this sentence is *isolated*. It does not hook up to any other statements of the theory; it does not help us derive new predictions that take advantage of the distinction between unholy water and regular water. Adding it to theory is like adding a piston that does not turn anything to an engine.
 - C. Perhaps surprisingly, sciences tolerate isolated sentences more than might have been thought. For this reason, it remains difficult to preserve science while banning metaphysics.
- III. Another way the logical positivists tried to avoid metaphysics involved refusing to take what theories seemed to say about unobservable reality too seriously. For the positivists, the job of theories is not to get the *world* right. It is to get *experience* right.
- A. Acupuncture provides a nice example. One can respect the highly reliable (at least within a certain domain) predictions that the theory makes and the cures it brings about, without taking the theory’s talk about energy channels and such fully seriously.
 - B. For the logical positivists, the connections among theoretical terms are crucial, but they are crucial for deriving observations, not for describing reality.
 1. Many statements in a scientific theory do not have to be true to be good. They are not attempts to describe the world but are, instead, inference tickets, saying that it is all right to infer this from that.
 2. They can still play a needed role in a theory’s ability to take observational inputs and generate true observational outputs. This is the *instrumental* conception of scientific theories.
 3. The point of a theory is not to make true statements that go beyond observation but to make true statements about patterns in experience.

Essential Reading:

Nagel, “Experimental Laws and Theories,” in Balashov and Rosenberg, *Philosophy of Science: Contemporary Readings*, pp. 132–140.

Hempel, “Empiricist Criteria of Cognitive Significance: Problems and Changes,” in Boyd, Gasper, and Trout, *The Philosophy of Science*, pp. 71–84.

Supplementary Reading:

Rosenberg, *Philosophy of Science: A Contemporary Introduction*, chapter 4.

Questions to Consider:

1. We're pretty sure that some counterfactual statements are true (for example, "If I were to flip this switch, the light would come on"). What makes this statement true? What is it about the way the world is that "governs" how things would go if the world had gone differently? Do more complicated counterfactual statements, such as "Had Hitler not invaded the Soviet Union, he would have defeated England," have straightforward (though perhaps unknowable) truth values? Why or why not?
2. Acupuncture seems to be a reasonably effective theory. Within its domain, it generates some true and surprising predictions, and it seems to be of genuine therapeutic value. But the theory behind these predictions looks rather peculiar, at least when judged from the standpoint of Western science (the theory involves pathways through which life energy flows, for instance). If the theory generates reliable predictions, should scientists care whether it fits well with other theories? Why or why not?

Lecture Eight

Holism

Scope: In this lecture, we confront an elephant that has been in the room with Popper and the positivists: the problem of *auxiliary hypotheses*. No statement can be shown to be true or false without relying on background assumptions. Consequently, empirical tests can, strictly speaking, show us only that something is wrong somewhere in our theory. This makes serious mischief for Popper's notion of a crucial test and for the positivists' program of establishing empirical meaning for individual sentences. Quine's *holism* is radical. He argues both that any statement can be preserved no matter how experience goes and that no statement is beyond the reach of revision on the basis of experience. Quine's hugely influential argument has been seen by many as an assault on the objectivity of science.

Outline

- I. A hypothesis such as "All copper conducts electricity" does not have any observational implications by itself—taken by itself, it is neither verifiable nor falsifiable.
 - A. Popper and the positivists understood this point, but they tended to underappreciate its philosophical significance.
 - B. We need some straightforward additional premises (for example, "This object is made of copper" and "This machine is built in such a way that the arrow will move to the right if an electric current is passing through it") in order to get an observable consequence, such as "The arrow will move to the right." These are called *auxiliary hypotheses*.
 - C. Strictly speaking, some rather peculiar auxiliary hypotheses are also needed (for example, "Electrical conductivity does not vary with the color of the experimenter's shirt").
 - D. An unexpected prediction shows only that at least one statement in our theory is false. Logic by itself will not tell us which statement(s) is (are) false.
 - E. This makes clear mischief for Popper's contention that science is distinguished by the way it tries to falsify its hypotheses. Experience and logic will not, without some help from us, falsify any given hypothesis.
 1. Generally, Popper does not think it appropriate to shift blame to an auxiliary hypothesis. A scientist should specify in advance which hypothesis will be rejected if an unexpected observation is made.
 2. But Popper does permit "blaming" an auxiliary hypothesis under certain conditions. The main requirement is that the auxiliary hypothesis can be independently tested.
 3. As we've seen in a number of contexts, there are worries about whether this standard is too restrictive and about whether it is too permissive.
 4. It is striking that Popper writes as if auxiliary hypotheses are testable in isolation. Popper knew that no hypothesis is testable in isolation, but he often ignored this fact.
 - F. The logical positivists also wrote as if hypotheses are testable in isolation. The explanation seems to be that they did not see a big problem here.
- II. W. V. Quine's "Two Dogmas of Empiricism," published in *The Philosophical Review* in 1951 and as a book chapter in 1953, is often considered the most important philosophical article of the century. In it, he draws radical implications from this idea that hypotheses are not testable in isolation.
 - A. Quine combined the idea that our theories face experience only as groups, not as single statements (*holism about theory testing*), with the positivists' notions about meaning (as, roughly, testability). Holism about testing, says Quine, implies holism about meaning.
 1. This means that statements do not have empirical significance in isolation. Theories, not statements, are the bearers of cognitive significance.
 2. This makes mischief for the logical positivists' project of distinguishing metaphysical from non-metaphysical statements. We can know the meaning of a scientific statement without having any clear idea of which observations would bear positively or negatively on it.

- B. Quine's most striking departure from the logical positivists is his claim that there is no interesting distinction between analytic statements (true by virtue of meaning) and synthetic statements (true by virtue of fact).
 - 1. Wasn't the distinction between a paradigmatically analytic sentence, such as "All bachelors are unmarried," and a paradigmatically synthetic statement, such as "The average American bachelor is 5 feet, 10 inches tall," pretty clear and impressive?
 - 2. Quine thought not. His main argument was that the analytic/synthetic distinction does no valuable philosophical or scientific work. Nothing turns on whether "Force equals mass times acceleration" is a definition or an empirical statement.
 - 3. For Quine, we should treat all beliefs as contingent and knowable only *a posteriori*. Any belief can be revised in the course of experience.
- C. Quine's view has a major consequence. Theory is always underdetermined by data. Observation never forces particular changes to a theory.
 - 1. No beliefs are insulated from the possibility of revision.
 - 2. Conversely, any statement can be maintained, no matter what experience says. If we are willing to make enough modifications to other parts of our theory, we will always be able to preserve a commitment to the truth of any statement.
- D. Quine's famous metaphor is that of a web of belief. Experience impinges on the edges, but there are always many ways of distributing that force through the web. It is possible to keep any local belief in place if you are willing to move enough stuff around it.
 - 1. Having done away with the analytic/synthetic and *a priori/a posteriori* distinctions, there are no sharp divisions within the web between philosophy and science or between science and metaphysics.
 - 2. Changes in the web of belief are to be guided by simplicity (minimize the number of basic laws and basic kinds of objects) and conservatism (preserve as much of the old theory as you can). These are pragmatic criteria.
 - 3. It is an open question whether these pragmatic criteria have any connection to truth.
- E. Quine argued that no matter how much information comes in, it does not force us to a unique theory. But Quine was no relativist, because he thought one should be constrained by simplicity and conservatism. He was far from thinking all theories or webs equal.
- F. Quine defended underdetermination by all possible data: There will always be more than one theory to fit the data, no matter how much evidence comes in.
 - 1. In actual science, the problem more often consists of finding one theory that fits the data reasonably well, not of choosing among many such theories.
 - 2. One way to explain this would be if there were additional constraints on the web, beyond those of deductive logic and beyond pragmatic constraints.
 - 3. If there were a scientific method that told how to update the web, that would explain why choices are so limited. But for Quine, those claims about method could only themselves be part of the web.

Essential Reading:

Quine, "Two Dogmas of Empiricism," in Curd and Cover, *Philosophy of Science: The Central Issues*, pp. 280–301.

Supplementary Reading:

Gillies, "The Duhem Thesis and the Quine Thesis," in Curd and Cover, *Philosophy of Science: The Central Issues*, pp. 302–319.

Laudan, "Demystifying Underdetermination," in Curd and Cover, *Philosophy of Science: The Central Issues*, pp. 320–353.

Questions to Consider:

1. When we look at the logic of scientific testing, the underdetermination of theory by data looks like a serious problem. But it almost never seems to arise in the “real world.” How would you explain this discrepancy between the logic and the history of science?
2. Use your imagination and some extreme cases to test some of Quine’s striking claims. Can you describe a web of belief in which it makes sense to maintain that the world is flat? Can you describe a web of belief in which it makes sense to give up $2 + 2 = 4$ or “No hummingbird is a sumo wrestler”?

Lecture Nine

Discovery and Justification

Scope: We turn now to issues of confirmation and evidence. To what extent can a methodology or logic of inquiry legitimately constrain one's web of belief? John Stuart Mill systematized a number of techniques deployed in earlier empiricist approaches to inquiry. Mill's methods are enormously valuable and are still very much with us, but they can seem both curiously ambitious and curiously naïve when judged by contemporary lights. On the one hand, their relentless empiricism carries with it a number of crucial limitations. On the other hand, at least as classically understood, Mill's methods try to *generate* the correct hypothesis. That's more than most contemporary methodologists think possible; they offer no theories for *finding* good hypotheses, only for *evaluating* them.

Outline

- I. Notions of evidence and justification have loomed large in the background of our discussions of demarcation and meaningfulness. We now turn directly to such topics, and we begin with a discussion of scientific method. In the most general sense, the study of scientific method is the study of whatever scientists do that helps account for the distinctive epistemic successes of science.
 - A. In principle, one could offer an entirely descriptive theory of scientific method; this would merely report on whatever methods scientists employ.
 - B. But in fact, just about any theory of scientific method is also normative: It describes methods that are supposed to work; it gives advice about what one should do and explains why.
 - C. The originators of the modern idea of a scientific method saw it as a kind of recipe for attaining knowledge. They disagreed about which recipe was the right one, but any recipe would have to share some crucial features.
 1. The method tells the inquirer how to discover and formulate the right answer or at least the right candidate answers. A scientific method is a method of *discovery*.
 2. The answers settled on were justified because they resulted from the application of the correct method. A scientific method is a method of *justification*.
 3. The method should be as close to mechanical as possible. A method is supposed to minimize the need for luck or genius.
 - D. In the 19th century, the idea caught on that hypotheses are first formed, then tested. But it is not the classic conception of a scientific method, which has inquirers read the right explanation/theory out of the data.
- II. John Stuart Mill defended a classical and strongly empiricist conception of method. Some time-honored empiricist methodological principles receive an influential formulation from Mill and figure centrally in his theory of method. They have come to be known as *Mill's Methods*.
 - A. Mill was an extreme empiricist, holding that all statements, even those of mathematics, should be testable by experience.
 - B. Mill's Methods are designed to take observations as input and to produce the right causal hypothesis as output.
 - C. Mill's Method of Agreement applies when two or more instances of the phenomenon under investigation share only one circumstance in common.
 1. The method then tells us to infer a causal connection between the circumstance and the phenomenon. For example, a number of patients all have cirrhosis of the liver, and they share the property of being heavy drinkers. The Method of Agreement directs us to infer that cirrhosis is due to heavy drinking.
 2. But what we observe is a correlation, not causation. For this reason, this method will not always reveal what sort of causal connection links cases.
 3. Furthermore, it can mistake coincidences for causes, and it assumes that similar effects are always produced by similar causes.
 4. What Mill has really established here is that any condition that is not always present when the phenomenon occurs cannot be necessary for the phenomenon.

- D. The Method of Difference applies when cases in which the phenomenon occurs and cases in which it does not share all circumstances except for one. This method has us infer that the circumstance in which the two cases differ is causally connected to the phenomenon under investigation.
 - 1. The Method of Difference has one noteworthy advantage over the Method of Agreement: It makes use of both negative and positive instances. It takes into account cases in which the phenomenon of interest fails to occur, as well as cases in which it occurs.
 - 2. As with the Method of Agreement, however, the causal relationships may be more complicated than the method can handle.
 - 3. Thus, this method really only allows us to show that if a condition occurs both where our phenomenon does and where it does not, then that condition cannot be sufficient for our phenomenon.
 - E. The Joint Method of Agreement and Difference combines the power of the preceding methods. We use the Method of Agreement to figure out what cannot be necessary for our phenomenon, and we use the Method of Difference to figure out what cannot be sufficient. We hope to be left with the condition that is necessary and sufficient.
 - 1. This method can be difficult to apply; we need the similarities and differences to line up very conveniently if the method is to be straightforwardly applicable.
 - 2. Despite its sophistication and complexity, this method still runs into problems with complicated cases of causality.
 - F. The Method of Concomitant Variations generalizes the Joint Method of Agreement and Difference. It comes into play when two or more phenomena co-vary positively or negatively, and it has us infer a causal connection between the phenomena.
 - 1. The discipline of statistics has greatly increased the power and reliability of relatively primitive methods like Concomitant Variations.
 - 2. As it stands, the method is vulnerable to complicated cases of causation, as when a correlation is mediated by a third variable.
 - G. The Method of Residues applies when we know what part of a phenomenon is due to the effect of certain causes and has us infer that the rest of the phenomenon is due to those causes that remain.
 - 1. This method also has its uses, but causation is, again, more complicated than the method allows.
 - 2. The method assumes (falsely) that causes are always additive. Cream gravy makes biscuits taste better. Jelly makes biscuits taste better. But gravy and jelly together make biscuits disgusting.
- III. Mill's Methods are enormously useful. But they can't lead unproblematically from observations to the correct causal hypothesis.
- A. We've seen that they have trouble handling causal complexity.
 - B. Mill's methods apply only if we have a list of all the circumstances that *might* be relevant to the phenomenon in question. But just about anything might be causally relevant.
 - C. Critics of Mill's empiricism insist that we must bring some kind of category scheme or theory of relevance to experience *before* we are in a position to learn from observations. They claim that Mill is trying to make observation do the work of theory.
 - D. Without some theory or hypothesis, Mill's critics suggest, we cannot so much as gather data that bears on the question at all. Do we discover hypotheses in the data or impose hypotheses on the data?
 - E. Mill's Methods allow no role for hypotheses that make reference to unobservable objects. This is a very significant limitation.
- IV. Popper and the logical positivists drew an important distinction between the *context of discovery* and the *context of justification*.
- A. In the context of discovery, they said that there was nothing worth calling a rational method. Worthwhile scientific hypotheses are generated through luck, hard work, or creative genius, not by applying a method.
 - B. There can, however, be a logic or method for testing hypotheses once they have been generated from whatever source.
 - C. Old-fashioned methods of discovery have been making something of a comeback recently, especially in artificial intelligence.

Essential Reading:

Hung, *The Nature of Science: Problems and Perspectives*, chapters 3 and 5.

Supplementary Reading:

Laudan, “Why Was the Logic of Discovery Abandoned?” in Brody and Grandy, *Readings in the Philosophy of Science*, pp. 409–416.

Curd, “The Logic of Discovery: An Analysis of Three Approaches,” in Brody and Grandy, *Readings in the Philosophy of Science*, pp. 417–430.

Questions to Consider:

1. If the number of hypotheses that could account for a given bit of data is more or less unlimited, how is it that human beings often seem to light on promising hypotheses fairly readily? How do we manage to narrow down the field so effectively?
2. Some have thought that the context of discovery should be governed not by logic but by economics. We should formulate and pursue (though we probably shouldn’t believe) hypotheses that can be tested easily and cheaply. What are the strengths and the weaknesses of such an approach?

Lecture Ten

Induction as Illegitimate

Scope: Any attempt to develop an inductive logic must go through or around Hume's skepticism about induction. Hume argues that you have no reason at all to believe that the Sun will come up tomorrow. This belief is caused (Pavlov-style) by experience, but it is not in the least justified by experience or by anything else. In this lecture, we wrestle with Hume's argument, then turn to Popper's dramatic response to it. He agrees with Hume, but he denies that science needs to rely on inductive inference at all. We develop and assess Popper's deductive conception of science and find that there is a significant price to be paid for disallowing induction.

Outline

- I. As we saw last time, Popper and the logical positivists were concerned with what they called the context of justification, not with the context of discovery.
 - A. They were interested in rational reconstructions of scientific reasoning, just as they had been interested in rational reconstructions of scientific theories. They were less interested in how theories are discovered or used than in the logical and evidential relations that hold within science.
 - B. Accordingly, they were interested in the logic of confirmation: What relationship must a theoretical statement bear to observation statements in order to receive evidential support from them?
 - C. It is important to note just how far beyond observation science routinely goes. Even a very simple statement, such as "All copper conducts electricity," vastly surpasses every observation that will ever be made.
 - D. A logic of confirmation won't guarantee that if our premises are true, our conclusion will be true. The fact that conclusions far outrun observational evidence for them guarantees that we are not going to find deductive proof here.
 - E. We're asking instead what premises must be like so that they provide a good or adequate reason for accepting a conclusion. A reason can be excellent without being conclusive.
- II. This lecture begins our discussion of inductive logic.
 - A. In saying this, we construe *induction* broadly. In the broad sense, *induction* simply contrasts with *deduction*. Induction encompasses all (rationally defensible) inferences that are not deductively valid.
 - B. There is also a narrower sense in which inductive inference forms just a subclass of inductive inferences in the broad sense, and we'll start with induction in this sense.
 1. In the narrow sense, inductions are "more-of-the-same" inferences.
 2. A classic such inference is induction to an instance: This licenses the inference from "All observed Xs have property P" to "The next X observed will have property P."
 3. Inductive generalizations, such as inferring that "All copper conducts electricity" on the basis of observations of conductive copper, are of greater scientific importance, because in science, we are more often interested in laws or patterns than in particular facts.
 - C. We should note that not all such inferences are justified. For example, if this is your first Teaching Company course, you would not infer from "All observed Teaching Company courses concern the philosophy of science" to "All Teaching Company courses concern the philosophy of science." But a large and varied sample of conductive copper does, we think, provide reason for thinking that all copper conducts electricity.
- III. David Hume offers a famous argument designed to show that inductive arguments are entirely unjustified.
 - A. For Hume, no number of observations of the Sun rising confers any evidential support for the conclusion that the Sun will rise tomorrow.
 - B. It is clear that science and common sense assume the legitimacy of some such inferences. If Hume's argument were to succeed, science would seem to be on an evidential par with superstitions, paranoid delusions, and so on.

- C. The argument first notes that no deductive justification of induction is possible.
 - 1. The fact that all observed pieces of copper have conducted electricity does not guarantee that the next piece of copper will do so (much less that all copper does so).
 - 2. We could have a valid deductive argument that the next piece of copper will conduct electricity if we could help ourselves to such a premise as “The future will resemble the past” or, less vaguely, “Future copper will resemble past copper with respect to conductivity.” But that premise is just what we are trying to establish; we cannot help ourselves to it.
 - D. Hume then argues that no inductive justification of induction is possible. Why should the fact that induction has worked well in the past count as a reason for thinking it will be reliable in the future?
 - 1. Whereas induction assumes that the future will be very much like the past, counterinduction predicts that the future will be unlike the past. For example, gamblers who have lost 10 hands in a row infer that they are due for things to get better.
 - 2. Counterinduction seems as if it could be justified in much the same way that we’re tempted to justify induction, namely, by appealing to its track record. But counterinduction seems utterly unjustified; thus, induction also appears utterly unjustified.
 - E. Hume did not think that we could or should refrain from performing inductive inferences. He would just have us realize that we are not governed by reason when we do so.
- IV. Popper accepted Hume’s argument but thought that induction played no role in science at all.
- A. For Popper, what matters is falsification, not confirmation, and scientific theories can be falsified using only observation and deductive logic. One black swan falsifies “All swans are white.”
 - B. In Popper’s view, scientists should not try to confirm their theories and, thus, do not need to reject Hume’s argument. The most we can ever say in favor of a theory or hypothesis is that it has survived strenuous attempts to falsify it. Popper called this *corroboration*. Corroboration does not indicate that a theory is healthy, only that it is not yet dead.
 - C. Popper denied that a theory’s corroboration is any predictor of future success. He had to deny it, because otherwise, he would have been relying on induction by arguing that the past survival of tests is evidence for the future survival of tests.
 - D. Popper’s view has trouble explaining why it is rational to prefer corroborated theories to untested theories.
 - 1. He seemed to say that the practice of science simply includes preferring corroborated theories. This is part of what makes science science. But that’s undermotivated; we’d like a reason to prefer the predictions of corroborated theories to those of untested theories.
 - 2. On Popper’s behalf, perhaps the best we can say is that we have no reason to drop a theory until it fails a test. But there could be lots of reasons to drop a theory if we don’t think it’s supported by the evidence.
 - 3. The problem looks even worse when we apply science to practical matters. Does it make sense to get on an airplane if one does not think past performance is any indicator at all of future performance?
 - E. Popper thought scientific theories aim at the truth, though in his opinion, they can never get any evidence that they have attained the truth.

Essential Reading:

Lipton, “Induction,” in Curd and Cover, *Philosophy of Science: The Central Issues*, pp. 412–425.

Popper, “The Problem of Induction,” in Curd and Cover, *Philosophy of Science: The Central Issues*, pp. 426–432.

Supplementary Reading:

Salmon, “Rational Prediction,” in Curd and Cover, *Philosophy of Science: The Central Issues*, pp. 433–444.

Questions to Consider:

1. Do you think that people sometimes make counterinductive inferences, or do you think that inferences that look counterinductive (for example, that a team with a terrible record is “due” for a win) are really inductive inferences in disguise?
2. Popperians (and others) sometimes claim that belief has no place in science; we might manifest belief in scientific results when we *use* them to build bridges and such, but science itself remains detached from belief and similar commitments. What are the strengths and weaknesses of such a conception of science?

Lecture Eleven

Some Solutions and a New Riddle

Scope: In this lecture, we consider and reject solutions to Hume's puzzle grounded in the law of large numbers and in the meaning of the term *rational*. We then turn to the *pragmatic vindication* of induction. It argues, not that induction will work or even that it is likely to work, but only that it will work if any other method will. On that basis, it is argued, we can rationally "bet on" induction. What kind and how much of a solution is this? We then turn to the work of Nelson Goodman, who offers a somewhat maddening new riddle of induction, according to which too many, rather than too few, inductive inferences appear justified.

Outline

- I. One might think that Hume's skepticism about induction runs afoul of a straightforward mathematical result.
 - A. The *law of large numbers* is a mathematical theorem. Informally presented, it states that, by taking a large enough random sample of a population, we can attain as high a probability as we would like of coming as close as we would like to knowing the frequency of a trait in the population.
 1. This law is often misunderstood. It does not require that we sample a high proportion of the population; it is the absolute size of the sample that matters. This is crucial to such activities as polling.
 2. A random sample is one in which each member of the sample has the same probability of being chosen.
 3. The law of large numbers seems to say that if we can get a suitably large sample, induction is just about guaranteed to work.
 - B. This response to Hume falters on the notion of randomness.
 1. We have no solid reason to believe our scientific samples to be random.
 2. We have some reasons to believe our samples are nonrandom. When we make a scientific claim about copper and electrical conductivity, it is not about early-21st-century copper on Earth. It is about all copper everywhere in the universe. Our experience looks tiny and nonrandom in the face of such considerations.
- II. The *ordinary language solution* to Hume's problem says that accepting some inductive arguments is part of what it means to be rational. Asking why it is rational to think that the Sun will come up tomorrow amounts to asking why it is rational to be rational. Induction is built into our notion of reason.
 - A. The ordinary language solution shows that induction cannot be justified by appealing to anything more fundamental than induction itself. Induction cannot be given a backward-looking justification; it cannot be derived from a more basic principle. Philosophers sometimes call this a *validation*.
 - B. But the ordinary language solution leaves intact the question of whether induction can get a kind of forward-looking justification, an explanation of what it is good for. Why, if at all, is induction a good way to get at the truth, to make predictions, and so on?
 - C. Although we also can't defend deduction without using deduction, there seems to be a big difference between induction and deduction.
 1. We can show to our satisfaction how deductive arguments serve the purpose of preserving truth and clarifying thoughts. Our rules preserve truth because they make explicit only what had been implicit in our premises.
 2. But we do not have an analogous understanding of why induction should do what it is supposed to do, namely, to extend our knowledge to unobserved cases. It can seem miraculous that we can go from a small sample to a grand conclusion.
- III. Our next solution, the *pragmatic vindication* of induction, perhaps wisely lowers its sights. It tries to show, not that induction will work, but that it will work if any method will.
 - A. The argument defends a simple version of inductive inference called *simple enumerative induction* (or the *straight rule*). This just says that we should infer that the entire population has a trait in whatever proportion that trait is exhibited in our sample.

- B. This argument does not assume that there is *a* proportion of the trait in the population as a whole. Maybe the proportion of copper in the universe that conducts electricity fluctuates wildly without ever settling on a value. If that turns out to be the case, no method can succeed, because there is no correct answer to the question.
 - C. But if there is a correct answer, a correct proportion of the trait in the sample, then an infinite application of the straight rule is guaranteed eventually to settle on that answer, and that cannot be said of any other method.
 - 1. The idea is that infinite sampling would eventually have to generate a random and, hence, representative sample.
 - 2. Other methods might get the right results faster, but if any method gets the right result, induction will eventually get there, too.
 - D. But we can show that there are still infinitely many rules that are guaranteed to work if any method will.
 - 1. All such methods appeal to an *a priori* component, a background belief about what the world is like that is not derived from the features of our sample. If there are three colors of marbles in an opaque jar, we might start with the idea that they each appear one third of the time.
 - 2. Such methods will work as well as the straight rule does provided that the background beliefs disappear as the sample size gets bigger.
 - 3. But this means we have not gotten anywhere, because having an infinite number of rules that make incompatible recommendations is a lot like having no defensible rule at all.
 - E. The pragmatic vindication is not dead yet, because there does seem to be a basis for preferring the straight rule.
 - 1. The other methods allow for different results without any change in the observations. If I differentiate between light-green marbles and dark-green marbles, there are now four colors of marbles in my jar. My method now gives me a different outcome, but I have changed only my language, not the data. If our choices about language determine our beliefs about the marbles, our beliefs seem arbitrary.
 - 2. Perhaps, then, the straight rule does have a special status, and we may have found a limited (since in the long run we're all dead) but significant defense of induction.
- IV. Nelson Goodman's "new riddle of induction" turns Hume's problem on its head. Goodman shows that our experience lends support to too many inferences of uniformity in nature, not too few. This problem dooms the pragmatic vindication.
- A. With his "grue" argument, Goodman claimed that even the straight rule allows for incompatible results, depending on the language one speaks.
 - 1. Call an object "grue" if it is first observed before January 1, 3000, and is green or if it is first observed after that time and is blue. There is no harm in introducing terms if they are clear.
 - 2. All emeralds ever observed have been grue; by the straight rule, then, we should expect emeralds first observed after January 1, 3000, to be blue. We are just projecting that the percentage of grue emeralds in the sample (vis-à-vis 100%) will match those in the population.
 - 3. At first blush, our evidence for the grueness of emeralds is every bit as good as our evidence for their greenness.
 - B. Goodman is not saying we should expect emeralds in the next millennium to be blue, any more than Hume was telling us to stop believing the Sun would rise. Both problems concern how good the *reasons* for our beliefs are.
 - C. It is far from clear that there's anything illegitimate about the term *grue*. It seems weird to us, but *green* would seem weird to us if we were "grue speakers."
 - D. There is no philosophical consensus on the notion of a real property that would include greenness but not grueness.
 - E. The same problem can be stated with unproblematic predicates and properties. The fact that all observed emeralds have the property of having been observed doesn't show that an emerald that won't be observed before January 1, 3000, has the property of having been observed.
 - F. Goodman took his riddle to show that the whole idea of an inductive logic is misguided. *Green* and *grue* bear the same logical relations to emeralds but aren't equally confirmed by observations of emeralds.

- V. Hume had us think that we could not find any real connections in nature. Goodman showed that connections or uniformities are too cheap to be valuable. The challenge is to figure out which connections or uniformities matter.

Essential Reading:

Achinstein, "The Grue Paradox," in Balashov and Rosenberg, *Philosophy of Science: Contemporary Readings*, pp. 307–320.

Hung, *The Nature of Science: Problems and Perspectives*, chapter 20.

Supplementary Reading:

Ladyman, *Understanding Philosophy of Science*, chapter 2.

Questions to Consider:

1. Can we believe (in God, in induction, etc.) on the basis of a pragmatic argument, or does belief respond only to evidential reasons?
2. Most defenses of induction focus on the long run, which can be very long indeed. To what extent do these defenses make it reasonable to use induction in the short run? If you are making only one bet on a roulette wheel that appears to be biased toward red, how does the (supposed) fact that red will turn up more in the long run affect what you should do here and now?

Lecture Twelve

Instances and Consequences

Scope: Carl Hempel offers a paradox that appears to be as frustrating as Goodman's. A black raven counts as a bit of evidence for "All ravens are black," right? Not so fast. This *instantial model* apparently implies that a white shirt supports the hypothesis that all ravens are black. As Goodman puts it, this opens up surprising prospects for indoor ornithology. We explore other problems with this account before turning to its enormously influential successor, the *hypothetico-deductive model* of confirmation. Though aspects of this approach seem indispensable, it, too, faces major challenges. Finally, we examine the idea of *inference to the best explanation* before leaving the topic of confirmation (for now). Can we make adequate sense of explanatory "betterness," and even if we can, is this a legitimate mode of inference?

Outline

- I. Let's back away from the problem of induction and just look at the notion of evidence itself. The positivists were looking for a logical relationship between an observation statement and a hypothesis such that the observation is evidence for the hypothesis. The most straightforward answer is provided by the *instantial model*, which says that an F that is G counts as evidence for "All Fs are G."
 - A. But Carl Hempel's *paradox of the ravens* seems to show that this model allows almost anything to count as evidence that all ravens are black.
 1. "All ravens are black" and "All non-black things are non-ravens" are logically equivalent. They are true under exactly the same conditions.
 2. It seems reasonable to insist that if a piece of evidence confirms a hypothesis, it also confirms any logically equivalent hypothesis.
 3. But, by this equivalent condition, any non-black non-raven is evidence for "All ravens are black." Thus, a white swan is evidence for "All ravens are black."
 - B. Hempel himself solved his problem by accepting that white swans provide evidence that all ravens are black. He denied that this is paradoxical: He said that it was a psychological illusion stemming from our mistaken sense that "All ravens are black" is only about ravens.
 1. Once we get over the mistaken impression that one hypothesis is "about" ravens and the other is "about" non-black objects (both are really about all objects in the universe), we can accept that a yellow pencil is evidence for "All ravens are black."
 2. We are letting background information about how many ravens there are in the universe compared to how many non-black things there are infect our intuitions, but that background information is not supposed to count in a logic of confirmation, because the relationship of evidence to theory is supposed to be formal.
 - C. A quite different approach to the raven paradox says that whether a piece of evidence confirms a hypothesis depends on such matters as how the information is collected.
 1. Evidence cannot confirm your hypothesis unless it is the kind of evidence that has a chance of falsifying (or at least disconfirming) it.
 2. If we discover that an object is yellow and then that it is a pencil (and, hence, not a raven), that observation does count in favor of our hypothesis because had the yellow object been a raven, our hypothesis would have been falsified or at least disconfirmed. But if we first learn that an object is a pencil and then that it is yellow, that observation has no bearing on our hypothesis.
 3. Hempel could not adopt an approach like this because he did not want background information or the order in which the information is received to matter to confirmation.
- II. For this reason, the raven paradox probably doesn't show that the *instantial model* is too weak. Surprisingly, the *instantial model* suffers, not from being too weak, but from being too strong.
 - A. As written, the model does not allow for confirmation of hypotheses that have any logical form other than "All Fs are G."

- B. Although we have granted that statements of this form are the most important ones for science, we would like our theory to allow us to get evidence for such statements as “There is at least one egg-laying marsupial,” a statement that is not of that logical form.
- C. More importantly, the instantial model applies only to statements that have observable instances. This is the main reason why the instantial model has been rejected.

III. The *hypothetico-deductive model* of confirmation is much more popular.

- A. According to this model, a hypothesis is confirmed by any evidence that the hypothesis entails. If my hypothesis says that the early bird gets the worm, then evidence that birds that hunt early weigh more than birds that do not hunt early counts in favor of my hypothesis.
- B. This model is free of the restrictions that plagued the instantial model. It allows us to say that the wave theory of light was confirmed when it was noticed that there is a bright spot in the middle of the shadow of a circular disk, even though we can’t directly observe light being a wave.
- C. But the hypothetico-deductive model allows a hypothesis to be confirmed by totally irrelevant data.
 1. Suppose my hypothesis is “Beagles weighing 2,000 pounds once roamed the Earth.” This hypothesis implies “Either 2,000-pound beagles once roamed the Earth or it is sunny today (or both).”
 2. Why does this implication hold? Because if the *if* part of the sentence is true, then the *then* part must be true.
 3. Suppose it is sunny today. That is enough to make “Either 2,000-pound beagles once roamed the Earth or it is sunny today” true. Because my original hypothesis implies a statement that was established as true by observation, my original hypothesis has been confirmed. Thus, a sunny day can count as evidence that 2,000-pound beagles once roamed the Earth.
 4. As we’ve seen before, any attempt to make room for apparently sensible cases tends to make room for apparently ridiculous cases.

IV. The model of *inference to the best explanation* requires that the hypothesis not merely entail the data but explain it.

- A. According to this model, a hypothesis is confirmed if the hypothesis would, assuming it to be true, provide the best explanation for the observed data.
- B. Sherlock Holmes used this approach quite a lot and misleadingly called it deduction. When Holmes inferred that the butler did it, he did so because that hypothesis does not just imply the facts; it (along with suitable auxiliary hypotheses) explains them.
- C. Like the hypothetico-deductive model, the inference to the best explanation model allows hypotheses about unobservables to receive evidential support. When a physicist sees a streak in a cloud chamber and says it is evidence for the presence of an electron, an inference to the best explanation is being performed.
- D. The name of this model is a bit misleading, given that sometimes, the best available explanation isn’t good enough. The detective may have a number of suspects, none of whom can legitimately be accused.
- E. Obviously, we will need a clearer understanding of scientific explanation than we currently possess if this model is to really work. But even apart from that problem, we can see that the notion of a better or best explanation is vexed.
 1. One notion of explanatory “betterness” would have us infer to the most plausible explanatory hypothesis. This is like saying that the team that scores the most points will win the game.
 2. The other main notion of explanatory “betterness” is loveliness, rather than likelihood. We should infer to the hypothesis that best accounts for the data or the hypothesis that would, if true, provide the greatest understanding of the data.
 3. But explanatory loveliness is a tricky notion. It can’t, for instance, amount to making the evidence maximally likely. On that view, if we draw a queen of hearts from a deck, we should infer that the entire deck consists of queens of hearts.
 4. Further, do we have any good reason for thinking that the hypothesis that makes the greatest contribution to our understanding is more likely to be true than a hypothesis that makes a smaller contribution?

Essential Reading:

Godfrey-Smith, *Theory and Reality: An Introduction to the Philosophy of Science*, chapter 3.

Hung, *The Nature of Science: Problems and Perspectives*, chapter 21.

Supplementary Reading:

Hempel, “Studies in the Logic of Confirmation,” in Brody and Grandy, *Readings in the Philosophy of Science*, pp. 258–279.

Harman, “The Inference to the Best Explanation,” in Brody and Grandy, *Readings in the Philosophy of Science*, pp. 323–328.

Questions to Consider:

1. How closely do the various models of confirmation we’ve looked at resemble what you think actually goes on in science or in ordinary life? Can your everyday inferences be reconstructed as inductive generalizations, inferences to the best explanation, or other models? What, if anything, gets left out of such reconstructions?
2. Do you think that simplicity, elegance, and explanatory loveliness are marks of truth? What reasons can you offer in support of your answer?

Timeline

| | |
|-------------------------------------|---|
| 6 th century B.C.E. | Thales asserts that water is the “primary principle.” This is arguably the first attempt at scientific explanation and at a scientific reduction. |
| 4 th century B.C.E. | Aristotle develops a systematic, sophisticated approach to scientific inquiry, involving both methodological and substantive advances. |
| c. 300 B.C.E. | Euclid develops the standard presentation of geometry, which stood as a model of scientific perfection for 2,000 years. |
| c. 400 C.E. | Evidence of sophisticated reasoning about probability appears in the Indian epic <i>Mahabharata</i> . |
| 1543 | Nicholas Copernicus puts forward the first detailed proposal that the Earth is a planet orbiting the Sun. The work was published with a preface by Andreas Osiander indicating that the theory should be treated as a calculating tool, not as a description of reality. |
| 1583–1632 | Galileo Galilei argues for the literal truth of the Copernican system, formulates a law of falling bodies and a law governing the motion of pendulums, applies the telescope to celestial phenomena, articulates a principle of the relativity of inertial motion, and generally develops a quantitative and observational approach to motion. |
| 1605–1627 | Francis Bacon develops the first systematic inductive method, a plan for attaining and increasing knowledge on the basis of experience. |
| 1609 | Johannes Kepler formulates his first two laws of planetary motion (the third law would have to wait 10 years). |
| 1628 | William Harvey establishes the circulation of the blood and the heart’s function as a pump. |
| 1633–1644 | Rene Descartes invents analytical geometry and develops his highly influential physics. |
| c. 1660 | The basic mathematics of probability takes shape in the work of Blaise Pascal, Christian Huygens, and others. |
| 1660 | The Royal Society of London for the Improving of Natural Knowledge is founded. Early members of the Royal Society include Robert Boyle, Christopher Wren, Robert Hooke, John Locke, and Isaac Newton. The Royal Society agitates in favor of experimental knowledge and against scholasticism and tradition. Many members are particularly interested in observational knowledge of witchcraft. |
| 1661–1662 | Robert Boyle takes major steps toward the separation of chemistry from alchemy, and he determines that the pressure and volume of a gas are inversely proportional. Boyle’s “corpuscularian” conception of matter greatly influenced John Locke. |
| 1666 | By this time, Isaac Newton had developed the fundamental principles of calculus, had formulated the principle of universal gravitation, and had established that white light consists of light of all colors of the spectrum. |
| 1673 | Moliere, in his play <i>Le malade imaginaire</i> , makes fun of the explanation that opium puts people to sleep because it has a “dormitive virtue.” |
| 1678 | Christian Huygens puts forward a version of the wave theory of light. |

| | |
|------------------|--|
| 1687 | Isaac Newton publishes his monumental <i>Principia</i> , which contains all the basic features of his mechanics, including his explicitly absolute conception of space and time. |
| 1690 | John Locke publishes his masterpiece, <i>An Essay Concerning Human Understanding</i> . |
| 1704 | Isaac Newton defends the particle (or corpuscular) theory of light in his <i>Opticks</i> . |
| 1709, 1714 | Gabriel Daniel Fahrenheit constructs an alcohol thermometer and, five years later, a mercury thermometer. |
| 1710 | Publication of George Berkeley's most important work, <i>A Treatise Concerning the Principles of Human Knowledge</i> . |
| 1715–1716 | Gottfried Wilhelm Leibniz develops a sophisticated relational account of space and time through a critique of Newton's work. |
| 1738 | Daniel Bernoulli publishes an early version of the kinetic theory of gases. |
| 1739–1740 | David Hume publishes his most important work, <i>A Treatise of Human Nature</i> . |
| 1750s | Carl von Linne (also known as Carolus Linneaus) launches the modern taxonomic system involving genera and species. |
| 1751 | Benjamin Franklin publishes <i>Experiments and Observations on Electricity</i> . |
| 1763 | Thomas Bayes's paper containing his famous theorem is presented to the Royal Society by Bayes's friend Richard Price. |
| 1769 | James Watt patents his steam engine. |
| 1770s | Joseph Priestly isolates a number of gases, including "dephlogisticated air," soon to be renamed <i>oxygen</i> by Antoine Lavoisier. |
| 1777 | Lavoisier performs the experiments that doom the phlogiston theory of combustion. |
| 1789 | Lavoisier establishes that mass is conserved in chemical reactions and formulates the modern distinction between chemical elements and compounds. |
| 1795 | James Hutton publishes <i>Theory of the Earth</i> , considered by many to be the founding document of the science of geology. |
| 1808 | John Dalton's <i>New System of Chemical Philosophy</i> propounds the atomic theory of chemistry. |
| 1809 | Jean-Baptiste Monet de Lamarck proposes the first really significant theory of evolution. Lamarck emphasizes the heritability of acquired characteristics. |
| 1818 | Simeon Poisson deduces from Augustin Fresnel's wave theory of light the apparently absurd consequence that a bright spot will appear at the center of the shadow of a circular object under certain conditions. Dominique Arago almost immediately verifies the prediction, however. |
| 1824 | Nicolas Léonard Sadi Carnot, despite relying on a conception of heat as a kind of substance, works out many of the central ideas of thermodynamics. |
| 1826 | Nikolai Ivanovich Lobachevsky produces a geometry that replaces Euclid's Fifth Postulate and allows more than one line parallel to a given line to pass through a fixed point. |
| 1830 | August Comte distinguishes theological, metaphysical, and positive stages of history, giving currency to the term <i>positivism</i> . |

| | |
|------------------|--|
| 1832 | Poisson proves a version of the law of large numbers and offers the clearest distinction yet drawn between “relative-frequency” and “degree-of-belief” approaches to probability. |
| 1840 | William Whewell develops a conception of scientific methodology that is hypothetical rather than purely inductive. In the same work, Whewell introduces the term <i>scientist</i> into the English language. |
| 1844 | Adolphe Quetelet argues that the bell-shaped curve that had been applied to games of chance and to astronomical errors could also apply to human behavior (for example, to the number of murders in France per year). |
| 1850 | Rudolph Julius Emanuel Clausius, generalizing Carnot’s work, introduces a version of the second law of thermodynamics. |
| 1859 | Charles Darwin publishes his epoch-making <i>On the Origin of Species</i> . |
| 1861 | James Clerk Maxwell reduces light to electromagnetic radiation. |
| 1866 | Gregor Mendel develops his theory of heredity involving dominant and recessive traits. |
| 1869 | Dmitri Ivanovich Mendeléev develops his periodic table of the elements. |
| 1870s | Ludwig Boltzmann offers two different reconciliations of the time directionality of the laws of thermodynamics with the time reversibility of the basic laws of motion. |
| 1878 | In Leipzig, Wilhelm Wundt establishes the first laboratory for physiological psychology. |
| 1879 | Gottlob Frege publishes his <i>Begriffsschrift</i> , arguably the founding document of modern mathematical logic. |
| 1887 | A. A. Michelson and E. W. Morley measure the speed of light as the same in all directions and thereby fail to detect any motion of the Earth with respect to the aether. |
| 1889, 1892 | G. F. Fitzgerald and H. Lorentz independently suggest that the null results of the Michelson-Morley experiments can be explained on the assumption that physical objects (such as measuring devices) contract at speeds approaching that of light. |
| 1892 | C. S. Peirce argues that there is no compelling scientific or philosophical reason for accepting determinism. |
| 1895 | X-rays are discovered by W. C. Röntgen. |
| 1900 | Max Planck introduces the “quantum theory,” according to which light and energy are absorbed and emitted only in bundles, rather than continuously. |
| 1902 | Ivan Pavlov carries out his well-known experiments involving learning and conditioned responses. |
| 1905 | Bertrand Russell publishes “On Denoting,” which becomes a paradigm of philosophical analysis. |
| 1905 | Albert Einstein publishes enormously important papers that, among other things, formulate the special theory of relativity and help explain Planck’s quantum theory. |
| 1912 | John Watson advocates behaviorism as the scientifically appropriate approach to psychology. |
| 1912 | A. L. Wegener proposes a unified theory of continental drift. |

| | |
|------------|--|
| 1913 | Niels Bohr publishes “On the Constitution of Atoms and Molecules,” which is often taken to contain the first theory of quantum mechanics. |
| 1915 | Einstein publishes his general theory of relativity. |
| 1919 | A team led by Arthur Eddington obtains experimental confirmation of Einstein’s hypothesis that starlight is bent by the gravitational pull of the Sun. |
| 1923 | Louis Victor de Broglie suggests that the wave-particle duality applies to matter as well as to light. |
| 1926 | Frank Ramsey, in “Truth and Probability,” lays much of the foundation for a rigorous interpretation of probabilities as degrees of belief. |
| 1926 | Max Born interprets electron waves probabilistically; the electron is more likely to be found in places where the square of the magnitude of the wave is large than where it is small. |
| 1927 | Werner Heisenberg denies that an electron simultaneously possesses a well-defined position and a well-defined momentum. |
| 1927 | Percy Bridgman’s “The Operational Character of Scientific Concepts” is published. |
| 1927 | Bohr and others formulate the Copenhagen interpretation of quantum mechanics. |
| 1929 | Edwin Hubble observes that all galaxies are moving away from one another. |
| 1929 | Rudolf Carnap, Otto Neurath, and Hans Hahn publish “The Vienna Circle: Its Scientific Outlook,” a manifesto of logical positivism. |
| 1931 | Sewall Wright argues that random genetic drift plays a significant role in evolution. |
| 1931 | Kurt Gödel’s incompleteness proof is published. This shows that any axiomatic system powerful enough to include arithmetic will imply at least one provably false consequence. |
| 1934 | Karl Popper publishes <i>The Logic of Scientific Discovery</i> . |
| 1936 | The first edition of <i>Language, Truth and Logic</i> , by A. J. Ayer, appears. |
| 1942 | Ernst Mayr publishes <i>Systematics and the Origin of Species</i> , a watershed work in biological classification. |
| 1942 | Julian Huxley publishes <i>Evolution: The Modern Synthesis</i> , which unified many aspects of biological research that had been achieved through the work of R. A. Fisher, Sewall Wright, J. B. S. Haldane, and others. |
| 1945 | Carl Hempel’s “Studies in the Logic of Confirmation,” which includes the raven paradox, appears in print. |
| 1948 | Hempel and Paul Oppenheim publish the first major statement of the covering-law theory of explanation. |
| 1951 | W. V. Quine publishes “Two Dogmas of Empiricism.” It appears in book form in 1953. |
| 1953 | James Watson and Francis Crick ascertain the chemical structure of DNA. |
| 1954 | Nelson Goodman’s <i>Fact, Fiction and Forecast</i> , which includes the classic statement of the new riddle of induction, appears. |

| | |
|-----------------|--|
| 1961 | Ernest Nagel's <i>The Structure of Science</i> , which presents a sophisticated positivist conception of science and includes a classic account of scientific reduction, is published. |
| 1962 | <i>The Structure of Scientific Revolutions</i> , by Thomas Kuhn, appears in print. |
| 1963 | J. J. C. Smart famously argues, in <i>Philosophy and Scientific Realism</i> , that there are no laws in biology. Smart's work is also sometimes taken to mark the resurgence of interest in scientific realism. |
| 1963 | Murray Gell-Mann and George Zweig independently arrive at the notion of quarks. Zweig treats them as tiny particles, while Gell-Mann thinks of them more as patterns than as objects. |
| 1969 | Quine's essay "Natural Kinds," a landmark of philosophical naturalism, is published. |
| 1970 | Saul Kripke presents the causal (also known as historical chain) theory of reference in lectures that would eventually be published as <i>Naming and Necessity</i> . |
| 1970–1971 | The most important papers outlining Imre Lakatos's methodology of scientific research programs appear. |
| 1972 | Stephen Jay Gould and Niles Eldredge argue that evolution largely proceeds in fits and starts, rather than gradually. This is known as the <i>punctuated equilibrium</i> approach to evolution. |
| 1974 | Michael Friedman publishes an influential account of explanation as unification. |
| 1974 | <i>The Structure of Scientific Theories</i> , a volume edited by Frederick Suppe, appears in print. The book contains classic presentations of the "received view" of scientific theories and of the then-new semantic conception of theories. |
| 1974 | <i>Knowledge and Social Imagery</i> , a classic work in the <i>strong program</i> in the sociology of knowledge, is published by David Bloor. |
| 1975 | Paul Feyerabend's <i>Against Method</i> appears in print. |
| 1977 | Larry Laudan's <i>Progress and Its Problems</i> , which includes a classic statement of the pessimistic induction argument against scientific realism, is published. |
| 1980 | Bas van Fraassen publishes <i>The Scientific Image</i> , which details both his constructive empiricism and his approach to explanation. |
| 1981 | Paul Churchland defends an influential version of eliminative materialism, the view that folk psychology is radically false and will be replaced. |
| 1982 | An Arkansas judge decides, in <i>McLean v. Arkansas Board of Education</i> , that creation-science does not count as science. The case included testimony about the problem of demarcation and has occasioned a great deal of discussion. |
| 1983 | David Armstrong's <i>What Is a Law of Nature?</i> awakens interest in non-regularity accounts of physical laws. |
| 1985 | Steven Shapin and Simon Schaffer's <i>Leviathan and the Air-Pump</i> , an important work in historical sociology of knowledge, is published. |
| 1988 | David Hull's <i>Science as a Process</i> , which examines such matters as the social structure and reward system of science, is published. |
| 1990 | Helen Longino publishes <i>Science as Social Knowledge</i> , a major work concerning social structure and objectivity. |

- 1990..... Philip Kitcher’s influential work on the division of cognitive labor appears in *The Journal of Philosophy*.
- 1996..... Alan Sokal’s parody of postmodernism, “Transgressing the Boundaries: Toward a Transformative Hermeneutics of Quantum Gravity,” is published in *Social Text*, and Sokal reveals his hoax in *Lingua Franca*. This period sees the height of the so-called Science Wars.

Philosophy of Science
Part II
Professor Jeffrey L. Kasser



THE TEACHING COMPANY ®

Jeffrey L. Kasser, Ph.D.

Teaching Assistant Professor, North Carolina State University

Jeff Kasser grew up in southern Georgia and in northwestern Florida. He received his B.A. from Rice University and his M.A. and Ph.D. from the University of Michigan (Ann Arbor). He enjoyed an unusually wide range of teaching opportunities as a graduate student, including teaching philosophy of science to Ph.D. students in Michigan's School of Nursing. Kasser was the first recipient of the John Dewey Award for Excellence in Undergraduate Education, given by the Department of Philosophy at Michigan. While completing his dissertation, he taught (briefly) at Wesleyan University. His first "real" job was at Colby College, where he taught 10 different courses, helped direct the Integrated Studies Program, and received the Charles Bassett Teaching Award in 2003.

Kasser's dissertation concerned Charles S. Peirce's conception of inquiry, and the classical pragmatism of Peirce and William James serves as the focus of much of his research. His essay "Peirce's Supposed Psychologism" won the 1998 essay prize of the Charles S. Peirce Society. He has also published essays on such topics as the ethics of belief and the nature and importance of truth. He is working (all too slowly!) on a number of projects at the intersection of epistemology, philosophy of science, and American pragmatism.

Kasser is married to another philosopher, Katie McShane, so he spends a good bit of time engaged in extracurricular argumentation. When he is not committing philosophy (and sometimes when he is), Kasser enjoys indulging his passion for jazz and blues. He would like to thank the many teachers and colleagues from whom he has learned about teaching philosophy, and he is especially grateful for the instruction in philosophy of science he has received from Baruch Brody, Richard Grandy, James Joyce, Larry Sklar, and Peter Railton. He has also benefited from discussing philosophy of science with Richard Schoonhoven, Daniel Cohen, John Carroll, and Doug Jesseph. His deepest gratitude, of course, goes to Katie McShane.

Table of Contents

Philosophy of Science

Part II

| | |
|---|----------|
| Professor Biography | i |
| Course Scope | 1 |
| Lecture Thirteen Kuhn and the Challenge of History | 3 |
| Lecture Fourteen Revolutions and Rationality | 5 |
| Lecture Fifteen Assessment of Kuhn | 8 |
| Lecture Sixteen For and Against Method | 10 |
| Lecture Seventeen Sociology, Postmodernism, and Science Wars | 13 |
| Lecture Eighteen (How) Does Science Explain? | 16 |
| Lecture Nineteen Putting the Cause Back in “Because” | 19 |
| Lecture Twenty Probability, Pragmatics, and Unification | 22 |
| Lecture Twenty-One Laws and Regularities | 25 |
| Lecture Twenty-Two Laws and Necessity | 28 |
| Lecture Twenty-Three Reduction and Progress | 31 |
| Lecture Twenty-Four Reduction and Physicalism | 34 |
| Timeline | Part I |
| Glossary | 37 |
| Biographical Notes | Part III |
| Bibliography | Part III |

Philosophy of Science

Scope:

With luck, we'll have informed and articulate opinions about philosophy and about science by the end of this course. We can't be terribly clear and rigorous prior to beginning our investigation, so it's good that we don't need to be. All we need is some confidence that there is something about science special enough to make it worth philosophizing about and some confidence that philosophy will have something valuable to tell us about science. The first assumption needs little defense; most of us, most of the time, place a distinctive trust in science. This is evidenced by our attitudes toward technology and by such notions as who counts as an expert witness or commentator. Yet we're at least dimly aware that history shows that many scientific theories (indeed, almost all of them, at least by one standard of counting) have been shown to be mistaken. Though it takes little argument to show that science repays reflection, it takes more to show that philosophy provides the right tools for reflecting on science. Does science need some kind of philosophical grounding? It seems to be doing fairly well without much help from us. At the other extreme, one might well think that science occupies the entire realm of "fact," leaving philosophy with nothing but "values" to think about (such as ethical issues surrounding cloning). Though the place of philosophy in a broadly scientific worldview will be one theme of the course, I offer a preliminary argument in the first lecture for a position between these extremes.

Although plenty of good philosophy of science was done prior to the 20th century, nearly all of today's philosophy of science is carried out in terms of a vocabulary and problematic inherited from logical positivism (also known as logical empiricism). Thus, our course will be, in certain straightforward respects, historical; it's about the rise and (partial, at least) fall of logical empiricism. But we can't proceed purely historically, largely because logical positivism, like most interesting philosophical views, can't easily be understood without frequent pauses for critical assessment. Accordingly, we will work through two stories about the origins, doctrines, and criticisms of the logical empiricist project. The first centers on notions of meaning and evidence and leads from the positivists through the work of Thomas Kuhn to various kinds of social constructivism and postmodernism. The second story begins from the notion of explanation and culminates in versions of naturalism and scientific realism. I freely grant that the separation of these stories is somewhat artificial, but each tale stands tolerably well on its own, and it will prove helpful to look at similar issues from distinct but complementary angles. These narratives are sketched in more detail in what follows.

We begin, not with logical positivism, but with a closely related issue originating in the same place and time, namely, early-20th-century Vienna. Karl Popper's provocative solution to the problem of distinguishing science from pseudoscience, according to which good scientific theories are *not* those that are highly confirmed by observational evidence, provides this starting point. Popper was trying to capture the difference he thought he saw between the work of Albert Einstein, on the one hand, and that of such thinkers as Sigmund Freud, on the other. In this way, his problem also serves to introduce us to the heady cultural mix from which our story begins.

Working our way to the positivists' solution to this problem of demarcation will require us to confront profound issues, raised and explored by John Locke, George Berkeley, and David Hume but made newly urgent by Einstein, about how sensory experience might constitute, enrich, and constrain our conceptual resources. For the positivists, science exhausts the realm of fact-stating discourse; attempts to state extra-scientific facts amount to metaphysical discourse, which is not so much false as meaningless. We watch them struggle to reconcile their empiricism, the doctrine (roughly) that all our evidence for factual claims comes from sense experience, with the idea that scientific theories, with all their references to quarks and similarly unobservable entities, are meaningful and (sometimes) well supported.

Kuhn's historically driven approach to philosophy of science offers an importantly different picture of the enterprise. The logical empiricists took themselves to be explicating the "rational core" of science, which they assumed fit reasonably well with actual scientific practice. Kuhn held that actual scientific work is, in some important sense, much less rational than the positivists realized; it is driven less by data and more by scientists' attachment to their theories than was traditionally thought. Kuhn suggests that science can only be understood "warts and all," and he thereby faces his own fundamental tension: Can an understanding of what is intellectually special about science be reconciled with an understanding of actual scientific practice? Kuhn's successors in sociology and philosophy wrestle (very differently) with this problem.

The laudable empiricism of the positivists also makes it difficult for them to make sense of causation, scientific explanation, laws of nature, and scientific progress. Each of these notions depends on a kind of connection or structure that is not present in experience. The positivists' struggle with these notions provides the occasion for our second narrative, which proceeds through new developments in meaning and toward scientific realism, a view that seems as commonsensical as empiricism but stands in a deep (though perhaps not irresolvable) tension with the latter position. Realism (roughly) asserts that scientific theories can and sometimes do provide an accurate picture of reality, including unobservable reality. Whereas constructivists appeal to the theory-dependence of observation to show that we help constitute reality, realists argue from similar premises to the conclusion that we can track an independent reality. Many realists unabashedly use science to defend science, and we examine the legitimacy of this naturalistic argumentative strategy. A scientific examination of science raises questions about the role of values in the scientific enterprise and how they might contribute to, as well as detract from, scientific decision-making. We close with a survey of contemporary application of probability and statistics to philosophical problems, followed by a sketch of some recent developments in the philosophy of physics, biology, and psychology.

In the last lecture, we finish bringing our two narratives together, and we bring some of our themes to bear on one another. We wrestle with the ways in which science simultaneously demands caution and requires boldness. We explore the tensions among the intellectual virtues internal to science, wonder at its apparent ability to balance these competing virtues, and ask how, if at all, it could do an even better job. And we think about how these lessons can be deployed in extra-scientific contexts. At the end of the day, this will turn out to have been a course in conceptual resource management.

Lecture Thirteen

Kuhn and the Challenge of History

Scope: Thomas Kuhn was more of a historian than a philosopher, but his 1962 book, *The Structure of Scientific Revolutions*, dealt logical positivism its mightiest single blow. It's not obvious how that could have happened—how exactly are his historical claims supposed to undercut the positivists' philosophical claims? In this lecture, we discuss the pattern Kuhn claims to find in the history of science—normal science punctuated by periods of revolution—and his explanation of this pattern via the notion of a *paradigm*. And we worry quite a lot about how the “ises” and the “oughts” of science bear on one another.

Outline

- I. The biggest blow to logical positivism came not from philosophy but from a historian of science, Thomas Kuhn. How exactly could historical claims bear on established philosophical doctrines?
 - A. The positivists and Karl Popper offered *rational reconstructions* of scientific reasoning, which tried to make the reasons behind the methods, decisions, and practices of science clear and explicit.
 - B. Such reconstructions do not attempt to provide empirically grounded descriptions of scientific behavior. They ignore many aspects of how science actually gets done. Popper and the positivists saw philosophy as an *a priori* discipline.
 - C. Nevertheless, such reconstructions should have some explanatory value. The fact that scientists follow a method or use a logic that the philosophers describe is supposed to be pivotal to the explanation of why science produces reliable results.
 - D. On the other hand, the underlying rationality of the scientific method(s) is not of much help in explaining various kinds of scientific failures and irrationalities.
 - E. For this reason, philosophers like the positivists made some assumptions about how science works, because they were confident that science as practiced exhibits rational method(s) for investigating nature better than any other undertaking does, and they assumed this fact was crucial to explaining the success of science.
- II. Kuhn insisted on mixing what the positivists had kept separate.
 - A. For Kuhn, the way to understand what is special about science is not to investigate an underlying method or logic but to look at all the mechanisms by which scientific views are adopted and modified. Science can only be understood “warts and all.”
 - B. Our best grip on such notions as scientific rationality comes from the history of science, not from the methodological principles of philosophers.
 - C. Kuhn was aware of the charge that he was confusing empirical disciplines with normative ones. Popper, for instance, agreed that much science was done as Kuhn described it but that only bad science was done that way.
 - D. Kuhn's view will be in trouble if his “warts-and-all” approach to science presents science as mostly warts.
- III. Kuhn held that science should be studied, in the first instance, by looking at what most scientists do most of the time. And he thought that historians, philosophers, and scientists had failed to understand *normal science*.
 - A. The sciences systematically misrepresent their history. They present it in a cumulative, triumphalist way. Kuhn went so far as to describe the history of science that is taught to scientists as a kind of brainwashing.
 - B. This approach, said Kuhn, has philosophical implications. The textbooks favor a broadly Popperian picture of science, full of heroes, bold conjectures, and dramatic experiments.
 - C. In fact, Kuhn argues, normal science is a relatively dogmatic and undramatic enterprise.
 - D. Normal science is governed by a *paradigm*.
 1. A paradigm is, first and foremost, an object of consensus.
 2. Exemplary illustrations of how scientific work is done are particularly important components of a paradigm. Scientific education is governed more by examples than by rules or methods.

- E. Paradigms generate a consensus about how work in the field should be done, and it is this consensus, not, as Popper thought, its perpetual openness to criticism, that distinguishes science from other endeavors.
- F. Normal science consists of puzzle-solving.
 - 1. The paradigm identifies puzzles, governs expectations, assures scientists that each puzzle has a solution, and provides standards for evaluating solutions.
 - 2. The paradigm is assumed to be correct. Normal science involves showing how nature can be fitted into the categories provided by the paradigm. Most of this work is detail-oriented.
 - 3. The paradigm tests scientists more than scientists test the paradigm. A failure to solve the puzzle reflects on the scientists' skills, not on the legitimacy of the problem.
- IV. But normal science has an important Popperian virtue: a remarkable power to undermine itself. A *crisis* occurs when a paradigm loses its grip on a scientific community.
 - A. Crises, according to Kuhn, result from *anomalies*—puzzles that have repeatedly resisted solution.
 - B. A crisis is a crisis of *confidence*; it is constituted by the reaction of the scientific community.
 - C. During such a crisis, the paradigm is subjected to testing and might be rejected.
 - D. Popper's mistake, according to Kuhn, is to have mistaken crisis science for normal science. Science could not achieve what it does if it were in crisis all the time.
 - E. Sometimes a new paradigm becomes ascendant. If this happens, a *scientific revolution* has taken place.
- V. How does Kuhn answer the charge that his normal science is bad science?
 - A. For Kuhn, dogmatism, crisis, and revolution are not failings of scientific rationality but enablers of scientific success.
 - B. Periods of crisis, sometimes followed by drastic rule changes, are crucial for inquiry, as long as they do not happen too frequently.

Essential Reading:

Kuhn, *The Structure of Scientific Revolutions*, chapters I–VIII.

Supplementary Reading:

Godfrey-Smith, *Theory and Reality: An Introduction to the Philosophy of Science*, chapter 5.

Bird, *Thomas Kuhn*, chapters 1–3.

Questions to Consider:

1. Do you think that normal science is as dogmatic as Kuhn says it is, as open-minded as Popper says it is, or somewhere in between?
2. How realistic a conception of the history of science was implicit in your scientific education? Were your science textbooks as simple-minded and triumphalist as Kuhn suggests that most science texts have been?

Lecture Fourteen

Revolutions and Rationality

Scope: This lecture examines Kuhn's (in)famously deflationary account of scientific rationality and progress across revolutions. Kuhn argues that proponents of competing paradigms will "see" different things in similar circumstances and, hence, that observation cannot adjudicate between paradigms. He insists that communication across paradigms will be partial at best and that rational discussion will be of limited use. He denies that we can make sense of science as getting closer to the truth. Nevertheless, Kuhn insists that he can make adequate sense of scientific progress and rationality. To what conclusion, exactly, do Kuhn's arguments lead? Has he really made science "a matter for mob psychology"?

Outline

- I. Though Kuhn's treatment of normal science is controversial, it is his treatment of scientific revolutions that has gotten people really worked up. Many thinkers find it deflating of science's aspirations and pretensions, because notions of rationality and truth play little role in Kuhn's explanation of the rise of a new paradigm.
 - A. A new paradigm will have achieved some impressive successes, but in general, it will be relatively undeveloped, and it will not be able to solve all the puzzles that the old paradigm could solve.
 - B. Often younger scientists, who are less invested in the old paradigm, switch to the new way of doing things. If their work looks promising enough, the new paradigm will continue to gain adherents, while proponents of the old paradigm die off.
 - C. But Kuhn rejected the triumphalist picture of old fuddy-duddies being superseded by clear-thinking young minds. Generational differences and other non-evidential factors come to the fore during a scientific revolution precisely because the evidence is inadequate to settle the matter.
 - D. In normal science, there is little room for the personal and idiosyncratic. In the freer conditions of crisis science, however, many personal factors can affect paradigm choice.
- II. Much of Kuhn's position can be summed up by his insistence that rival paradigms cannot be judged on a common scale. They are *incommensurable*. This means they cannot be compared via a neutral or objectively correct measure.
 - A. Standards of evaluation vary too much across paradigms to be of decisive use.
 1. Certain values are more or less permanent parts of science: predictive accuracy, consistency, broad scope, simplicity, and fruitfulness.
 2. But these values can be interpreted, weighed, and applied in different ways. They often conflict with one another.
 3. Thus, work in each paradigm is governed by scientific values, but each paradigm will hold work to the standards provided by that paradigm.
 4. Even within a paradigm, these values do not function as explicit principles but, rather, as shared habits and ways of seeing things. This is crucial for the proper function of science, but it limits the role of explicit, reasoned comparison of paradigms.
 - B. Effective communication across paradigms is very difficult.
 1. Like W. V. Quine, Kuhn adopts a holistic conception of meaning. Both are influenced by the positivists' idea that terms and statements get their meaning from their role in deriving observational consequences.
 2. Because the meaning of a term or statement derives from the role it plays in a theory, changes elsewhere in the theory or paradigm can bring about significant changes in the meaning of a term or statement.
 3. For this reason, Kuhn denies that a term such as *mass* means the same thing in Einstein's theory that it does in Newton's. Einstein offers a theory about different stuff, rather than an improved theory of the same stuff.
 4. For reasons such as these, proponents of different paradigms tend to talk past each other.

- C. Paradigm-neutral observations cannot be used to adjudicate between paradigms.
 - 1. For Kuhn, observation is *theory-laden*. What people see depends, in pertinent part, on what they already believe or expect. Seeing is less passive, less receptive than many had thought.
 - 2. Kuhn thus denies that we have access to a realm of observational evidence that is largely independent of theory and could, then, count as a source of meaning and evidence.
 - 3. Kuhn commits himself to rather extreme-sounding versions of this point. He says that, in an important sense, followers of different paradigms inhabit different worlds.
- D. Consequently, changing paradigms is, to some extent, like having a conversion experience. Because individual psychology is crucial to understanding why individuals change paradigms and because the senses of crisis and resolution are largely social phenomena, it is not hard to see why the Hungarian philosopher Imre Lakatos called Kuhn's picture one of mob psychology.

III. Science, for Kuhn, cannot be seen as straightforwardly cumulative, progressive, or truth-tracking.

- A. The history of science does not support a claim of progress. Einstein's physics resembles that of Descartes more than that of Newton in some key respects.
- B. Given that the victors write history, science is taught in a way that makes it seem more cumulative and progressive than it really is.

IV. On the other hand, Kuhn often wrote as if science does manifest a genuine tendency toward increasing problem-solving ability.

- A. Dogmatism and idiosyncrasy, for Kuhn, function in a complex social arrangement to produce desirable outcomes, just as in Adam Smith's economic model, individual selfishness produces socially desirable outcomes.
- B. It is unclear how Kuhn's trust and claim of progress can be reconciled with his arguments for incommensurability. Those discussions suggest that new paradigms solve different problems, not more or better problems.
 - 1. It is reasonably clear that Kuhn was not a complete relativist about science: He thought it the best method of investigating the natural world because it is good at generating and solving puzzles about nature.
 - 2. It is equally clear that Kuhn rejects the claim that science progresses in the sense of getting closer to the truth. Truth, for Kuhn, makes sense within paradigms but is unclear and dangerous when applied across paradigms.
 - 3. Kuhn sometimes goes so far as to deny the intelligibility of such notions as extra-paradigmatic truth or reality.

Essential Reading:

Kuhn, *The Structure of Scientific Revolutions*, chapters IX–XIII, plus the postscript.

Kuhn, "Objectivity, Value Judgment and Theory Choice," in Curd and Cover, *Philosophy of Science: The Central Issues*, pp. 102–118.

Supplementary Reading:

Godfrey-Smith, *Theory and Reality: An Introduction to the Philosophy of Science*, chapter 6.

Bird, *Thomas Kuhn*, chapters 4–5.

Questions to Consider:

1. How apt do you find the analogy between changing paradigms and undergoing a religious conversion? Insofar as the comparison is apt, how troubling should it be to scientists?
2. Do you think that science progressively gets closer to the truth? What evidence bears on this question? Do you think that science accumulates problem-solving ability? What evidence bears on this question?

Lecture Fifteen

Assessment of Kuhn

Scope: Kuhn's powerful and wide-ranging work demands that we ask questions of several different types: How accurate is his portrayal of patterns in science? To the extent that it is accurate, how acceptable is Kuhn's explanation of this pattern? Are his claims about perception psychologically and philosophically defensible? How philosophically sophisticated are his views of language and truth? We will discover that critics who object to Kuhn's radicalism and those who object to his traditionalism could have a surprising amount in common. Much of Kuhn's apparent radicalism derives from assumptions he shares with his empiricist predecessors.

Outline

- I. Kuhn has compelled philosophers to pay more careful attention to the history of science. But some have found Kuhn's descriptions and explanations of scientific episodes unconvincing.
 - A. Does normal science work as Kuhn said it did?
 1. Are scientists as committed to their paradigms as Kuhn suggested, or is there room for more Popperian detachment than Kuhn allowed?
 2. Are the contexts of discovery and justification as intertwined as Kuhn suggested, or are the guiding and justifying roles of paradigms more distinct than he realized?
 3. Are the elements of a paradigm as inseparable from one another as Kuhn believed?
 - B. Are normal science and revolutionary science as distinct from each other as Kuhn suggested?
 1. Some episodes of revolutionary science do not appear to have been preceded by crises.
 2. Some work that had revolutionary consequences required little or no change in previous beliefs.
- II. Kuhn's claims about incommensurability have attracted a great deal of largely unfavorable attention from philosophers. To what extent can Kuhn fairly be charged with making science a matter of "mob psychology"?
 - A. Are scientific values (or rules or methods) as incapable of adjudicating between paradigms as Kuhn claims?
 1. Kuhn does not have much to say about why science values such things as simplicity and explanatory power.
 2. If one can link such values to truth or similar epistemic goals, then some episodes in the history of science that look like matters of taste to Kuhn can be reconstructed as instances of rational theory choice.
 - B. How fraught with difficulty is communication across paradigms?
 1. It is not clear that we find the evidence of miscommunication and misunderstanding across paradigms that we should expect to find if Kuhn were right.
 2. It is not clear that we want to grant that meaning is as holistic as Kuhn says it is. If we consider fewer of a term's inferential connections essential to its meaning, then meanings can sometimes remain constant across paradigm shifts.
 3. Even if we grant that meaning is as sensitive to changes within a theory as Kuhn says it is, do such semantic changes generate the level of misunderstanding that Kuhn sometimes suggests they do?
 - C. Kuhn's rejection of paradigm-neutral observations has probably generated more criticism than any other aspect of his view.
 1. The influence of theory on observation is not all that powerful: The Sun still appears to rise, even after we learn that it does not.
 2. Kuhn tends to run together descriptions of visual experiences and the visual experiences themselves. Even if we grant that perception is significantly theory-laden, we need to leave ourselves room to say that nobody has ever seen the Sun move around the Earth, because there is no such state of affairs to be seen.

3. It is difficult to make clear sense of some of Kuhn's provocative comments about world change. When Kuhn claims that scientists inhabit different worlds, he needs to mean more than that they believe different things, but he must also avoid having scientists see what isn't there.
4. Kuhn insists that it is not just experience, concepts, and beliefs that change across paradigms; the world itself changes. This is to deny any use for a phrase such as "the real world."
- D. Perhaps the most important thing to note about this issue is that observations that are couched in a theory's terms do not thereby lose any ability to falsify that theory. The notion of a pre-Cambrian rabbit is stated in terms of a standard geological/biological theory. But that wouldn't prevent an observation from falsifying the theory. From the fact that our theories influence our perceptions, it doesn't follow that we can see only what our theories say is there.

III. How persuasive is Kuhn's skepticism about scientific truth?

- A. Kuhn inherits certain assumptions about what real knowledge would be from the logical positivists.
- B. He realizes that knowledge is messier than the positivists had thought. Observation and definition don't yield knowledge as straightforwardly as we might have hoped they would. For this reason, Kuhn backs away from talk of knowledge and truth.
- C. Arguably, what's needed is a different model of knowledge. In such a picture, a theory won't be understood as an impediment between oneself and the world. It will be thought of more as an investigative tool, one that allows us to build on and extend observational evidence.

Essential Reading:

McMullin, "Rationality and Paradigm Change in Science," in Curd and Cover, *Philosophy of Science: The Central Issues*, pp. 119–138.

Laudan, "Dissecting the Holist Picture of Scientific Change," in Curd and Cover, *Philosophy of Science: The Central Issues*, pp. 139–169.

Supplementary Reading:

Bird, *Thomas Kuhn*, chapters 6–7.

Nickles, ed., *Thomas Kuhn*.

Questions to Consider:

1. Discussions of Kuhn often contrast judgments of taste with rule-governed judgments of rationality. How impressed are you by this contrast?
2. How much commensurability do you think is needed for rational choice? When people choose between radically different options (maintaining a relationship versus accepting a job offer, joining the Peace Corps versus going to law school, and so on), to what extent do they represent these options on a common scale (for example, happiness)? To what extent does our inability to find a common scale limit our ability to make rational decisions?

Lecture Sixteen

For and Against Method

Scope: Imre Lakatos provides the first major attempt to reconcile much of the rationalism of the received view with Kuhn's historicism. His *methodology of scientific research programs* tries to accommodate both Popperian openness to criticism and Kuhnian attachment to theories. Methodological rules assess research programs in historical terms as *progressive* or *degenerating*. Paul Feyerabend, philosophy of science's great gadfly, sees Kuhn as glorifying dull, mindless scientific activity. In arguments alternately sober and outlandish, Feyerabend defends scientific creativity and "epistemological anarchism."

Outline

- I. Imre Lakatos put forward the first major post-Kuhnian theory of scientific methodology. Lakatos sought to reconcile Kuhn's historical approach to the philosophy of science with a much more robust role for scientific rationality.
 - A. Lakatos refused to share Kuhn's confidence in the actual practice of science. Having fought the Nazis during World War II and having been imprisoned for "revisionism" by the Hungarian government in the 1950s, Lakatos rejected Kuhn's notion that there is no higher scientific standard than the assent of the relevant community. Lakatos insisted on placing trust only in methods and rules, not in people or social practices.
 - B. Following Kuhn, Lakatos insisted that philosophical views about science had to be tested against the history of science. But he followed Popper and the logical positivists in thinking that a universal method survives the test of history.
- II. Lakatos's view is called *methodology of scientific research programs*. It can be seen as a compromise between the Popperian and Kuhnian approaches.
 - A. A research program is, for the most part, very like a Kuhnian paradigm.
 1. A research program includes a *hard core* of principles. This core is taken to be beyond criticism. Newton's three laws of motion and his law of gravitation form the hard core of Newtonian physics.
 2. A research program also includes a *protective belt* of claims that can be modified as needed to insulate the core from falsification.
 3. The protective belt permits a research program to develop over time. For Lakatos, a research program can be evaluated only over time, not at a time. Research programs constantly face anomalies but need not be rejected on that basis.
 4. A major difference between Lakatos's research programs and Kuhn's paradigms is that Lakatos permitted competing research programs to flourish at the same time.
 - B. Lakatos thought that research programs could be evaluated in an objective way by comparing them over time. He borrows a good bit from Popper here.
 1. A *progressive* research program modifies its protective belt in ways that generate new predictions. It generates its own research momentum.
 2. A *stagnant* or *degenerating* research program merely reacts to anomalies; it does not cope with them in ways that generate new predictions.
 3. Perhaps surprisingly, Lakatos puts less weight on the empirical correctness of the research program's predictions than he does on the program's ability to integrate problems smoothly into a progressive research agenda.
 4. Lakatos defends objective standards of evaluation but has only modest things to say by way of advice. We can know whether a research program is a good one only after the fact.
 5. It is not a rule of scientific rationality that one should abandon degenerating research programs for progressive ones. Philosophy of science cannot provide such advice; one might have reason, for instance, to think that the program will become progressive again.

- C. Lakatos argues that a philosophy of science is to be judged by how rational it makes the history of science look.
1. The history of science provides the data, and a philosophical research program is judged by how progressively it handles the data over time.
 2. Because a philosophical research program is supposed to make the history of science seem rational, philosophical history of science is supposed to be “Whiggish,” that is, written from a contemporary point of view. Lakatos takes it as a given that it was rational for scientists to reject Newton in favor of Einstein; the philosopher is supposed to explain why.
 3. Lakatos’s approach involves a great deal of rational reconstruction; philosophical histories aren’t supposed to be especially empirically accurate. Philosophers should write the history of science as their methodologies say it should have been.
 4. For Lakatos, the point of the history is logical, not empirical. The more problems the theory sets for itself that it knows how to approach, the more progressive the program looks. The more that other factors have to be called upon to explain scientific behavior, the more degenerative the program looks.
- III. Paul Feyerabend argues against any version of a scientific methodology. If you insist on having a rule governing scientific practice, only one will do: “Anything goes.”
- A. Feyerabend likes to make fun of other philosophers, and he doesn’t always accept his own arguments; sometimes their purpose is to “show how easy it is to lead people by the nose in a rational way.”
- B. Feyerabend’s most influential argument derives from historical cases and is, in that sense, recognizably Kuhnian in spirit.
1. Any set of rules, said Feyerabend, would, if followed, have prevented at least one important scientific advance.
 2. His central example concerns Galileo’s arguments for the Copernican hypothesis. Galileo’s genius involved overcoming observation, not following it, according to Feyerabend (because, for instance, a stone dropped from a tower should land away from the tower if the Earth is spinning).
 3. Galileo was also opposed by a massively supported theory; the whole Aristotelian approach to physics stood against Copernicus.
 4. In overcoming these formidable obstacles, says Feyerabend, Galileo used propaganda, unfair rhetoric, and intentionally bad arguments in the service of his worldview, and Feyerabend thought it a good thing that Galileo had done so.
- C. Whereas Kuhn deemphasizes methodological principles because he trusts the social practice of science, Feyerabend does so because he trusts and, indeed, celebrates individual creative scientific geniuses. He sees Kuhn as valorizing scientific drudgery.
- D. Feyerabend tries to link his celebration of scientific creativity to more traditional concerns, such as testability and evidence. Like Popper and Lakatos, Feyerabend thinks that theories could and should be tested against one another, rather than just against the world (or experience).
1. Because we want our theories to receive severe tests, we should develop and defend as great a variety of theories as possible. In order to maximize testing, we should struggle not to be limited by our sense of the plausible.
 2. Feyerabend is not much concerned with the “white noise” problem. His approach would generate lots and lots of theories but gives us little guidance about how to distribute our attention and resources among all these theories.
- E. Feyerabend is not, as he is sometimes taken to be, anti-science. Galileo and similar scientists are great heroes of his. But he believed that modern science resembles the Catholic Church of Galileo’s day: It stifles the spirit and imagination of those involved in it and bullies those who do not understand it. The scientific monopoly on legitimate intellectual authority, he believes, makes it a threat to democracy.

Essential Reading:

Godfrey-Smith, *Theory and Reality: An Introduction to the Philosophy of Science*, chapter 7.

Supplementary Reading:

Feyerabend, *Against Method*.

Larvor, *Lakatos: An Introduction*.

Questions to Consider:

1. Lakatos's approach to the history of science is unabashedly Whiggish. To what extent is Whig history appropriate in science or in philosophy? On the one hand, we're interested in reasons, not just in causes, when we look at science or philosophy. On the other hand, how can Lakatos be practicing history when he represents events as much more rational than they actually were?
2. We've seen most of Kuhn's defense against the charge that he's an epistemological anarchist. How do you think he would respond to Feyerabend's charge that Kuhn is a defender of drudgery?

Lecture Seventeen

Sociology, Postmodernism, and Science Wars

Scope: In the Kuhnian aftermath, sociology of science set itself up as a “successor discipline” to philosophy of science. The *strong program* in the sociology of science insists that beliefs should receive the same sort of justification, whether we think them true or false, well- or ill-founded. In particular, strong programmers maintain that decisions about which scientific theories to accept are determined by needs and interests (including, especially, social and political interests) of those making the decision. The notion of the *social construction of reality* receives careful attention. We also examine the highly controversial application of postmodernism to science, which prompted the physicist Alan Sokal’s successful submission of a parody essay to the journal *Social Text*.

Outline

- I. Nobody denies that social factors have some bearing on how science gets done. Social priorities affect which diseases are studied, for instance. But traditionally, social factors play a role only in setting questions, not in answering them. Kuhn blurs the line between the social and the evidential. A scientific crisis, for instance, is more a matter of confidence than of evidence, according to Kuhn.
- II. In the Kuhnian aftermath, a new approach to science emerged in the discipline of sociology that made much more of social factors and much less of epistemic ones than Kuhn had. The most influential version of this new approach was the *strong program* in the sociology of science, which emerged at the University of Edinburgh in the 1970s.
 - A. This new discipline set itself up as a “science of science,” a successor to the philosophy of science, which it regarded as misguided.
 - B. The centerpiece of the strong program is the *symmetry principle*, which requires that unreasonable or untrue beliefs (by our lights) receive the same kinds of explanation as reasonable or true beliefs.
 1. Strong programmers take a kind of anthropological look at the scientific community, its social norms, its structures of prestige and authority, and its practices for settling disagreements, without suggesting that any of these norms, structures, or practices are especially rational or truth-conducive.
 2. Beliefs are to be explained by local norms and non-epistemic interests. Strong programmers think that such notions as truth and rationality are unsuitable for scientific purposes. They do allow a role for notions about what a community considers true or rational. The scientific community thinks its beliefs and practices especially rational, but so do lots of other communities.
 3. Sometimes the kind of explanation at issue is not entirely clear. The most natural way to interpret some of the explanations is as causal hypotheses, but social and political interests rarely straightforwardly *determine* scientific opinions. A sociologist might make a convincing argument that certain scientific ideas would benefit a certain group, but that is not to say that the benefit explains why the views were adopted.
 - C. Particular works in the sociology of science are often illuminating, but the strong program is a program, not a particular claim, and it is the programmatic statement that it is never appropriate to explain beliefs in terms of truth, rationality, or evidence that has exercised philosophers.
 1. Sociologists think, for example, that evidence is more or less powerless to choose among theories—as we saw with Quine, too many theories can be compatible with the evidence. But sociologists seem to think that interests can sort through this underdetermination. Philosophers want to know why it is not just as unclear which theory best fits certain interests as it is which best fits the evidence.
 2. The strong programmers arguably share with (some of) the positivists an excessively narrow conception of evidence and reasoning. The more untainted by theory an observation would have to be in order to count as evidence, the easier it is to minimize the role of evidence in science. The more formal and rule-governed reasoning would have to be to count as reasoning, the smaller the role for reasoning in science. Kuhn can be unclear on these issues, but he at least did not contrast the realms of the social and the rational to the extent that his predecessors and successors did.

3. Kuhn flirted with relativism; the strong programmers adopt relativism. Any belief about the superiority of science or any other practice is to be explained from within local norms, and no such judgments have any standing outside such norms.
 4. Like all relativists, proponents of the strong program face a problem of self-reference. For the most part, strong programmers grant that their own views are to be explained in terms of the norms and interests governing their community, not in terms of accuracy. This presents something of a problem.
- D. Steven Shapin and Simon Schaffer's *Leviathan and the Air-Pump* provides an impressive illustration of the strong program in action. We've looked at philosophers' epistemological objections to the strong program, and this work allows us to consider some metaphysical objections.
1. Shapin and Schaffer study the rise of experimentation in England in the late 17th century. They locate a social function for experimentation: It was designed to settle disputes publicly and cooperatively. They suggest that the motivation for this was as much political as epistemic; in a time of religious wars, a method was needed for settling questions amicably.
 2. Shapin and Schaffer go so far as to suggest that such experimentalists as Robert Boyle were engaged in the manufacture of facts. They write, "It is ourselves and not reality that is responsible for what we know."
 3. This kind of language invites confusion, and what Shapin and Schaffer say seems to me misleading at best. Views such as this, according to which reality is made rather than found, are called *social constructivist*.
 4. Though people often suggest otherwise, being socially constructed does not imply being less than fully real. Buicks are socially constructed—the result of a complex social practice—yet they are thoroughly real.
 5. Something is real if its being a certain way doesn't depend on anybody's thinking that it *is* that way (this conception is due to C. S. Peirce). We need a conception of reality like this one to sort through the many confusions in this field.
 6. There are many different ways in which a term such as *social construction* could be used. Nations or corporations are real—their existence does not depend on what anyone in particular thinks about them—but they are also recognizably socially dependent in a way that Buicks are not. A decree can dissolve a company but not a Buick.
 7. The term *social construction* is most helpfully applied to things that are generally thought to have more independence from our practices than they actually do. Race is biologically unreal but socially real.

III. Postmodern approaches to science bear distant affinities to sociological approaches. Postmodernism comes out of the humanities and rests on very general claims about language and reality. Speaking somewhat loosely, philosophical postmodernism questions the ability of linguistic and other signs to represent anything worth calling real.

- A. Science's apparent success at "getting the world right" needed to be debunked given postmodernism's sense that the very notion of getting the world right is deeply flawed.
- B. For postmodernists, science is essentially a literary genre and nature, essentially a text. Postmodernists have drawn useful attention to rhetorical strategies and figurative language in science, but most people are unpersuaded by the idea that science, at the end of the day, consists of a rather tedious literary genre.
- C. Scientists have been unimpressed by the "one-size-fits-all" nature of most postmodernist criticism of science, while postmodernists have often thought scientists epistemologically and politically naïve and conservative.
- D. The stage was thus set for some brief but well publicized Science Wars in the 1990s, highlighted by the successful submission of a physicist's parody to a postmodern journal of science.
- E. The Science Wars generated more heat than light. It was inappropriate for the postmodernists to be as dismissive as they were of science, but it was also inappropriate for science's self-appointed defenders to treat science as above reproach or criticism.

Essential Reading:

Godfrey-Smith, *Theory and Reality: An Introduction to the Philosophy of Science*, chapters 8–9.

Supplementary Reading:

Bloor, “The Strong Programme in the Sociology of Knowledge,” in Balashov and Rosenberg, *Philosophy of Science: Contemporary Readings*.

Shapin and Schaffer, *Leviathan and the Air-Pump: Hobbes, Boyle, and the Experimental Life*.

Questions to Consider:

1. What do you think can be learned from literary approaches to science? What narrative and rhetorical features do you think loom large in scientific discourse, and what significance do these features have?
2. To what extent do you adopt something like the symmetry principle when you are explaining the actions of other people?

Lecture Eighteen

(How) Does Science Explain?

Scope: At the midway point of the course and with the radical wing of Kuhn's followers having gone down something of a dead end (or so it would seem to most philosophers), we return to logical empiricism to explore some ideas that have come to the philosophical fore in the time since the Kuhnian revolution. Many empiricists denied that science explains phenomena. The demand for explanation, they argued, leads inexorably to metaphysics—experience tells us only that something happens, not why it happens. Yet it seems that science does and should offer answers to “why” questions. Carl Hempel's *covering-law model* of explanation manages the delicate task of respecting empiricist scruples while forging genuine explanatory relations. Explanations are arguments telling us what to expect given the laws of nature. But does this attractive approach to explanation exclude legitimate but non-law-governed explanations from biology and the human sciences?

Outline

- I. Though explanation seems a central ambition of science, thinkers in the empiricist tradition have been somewhat suspicious of the notion of explanation.
 - A. It seems obvious that science tries to tell us not just what happens but why it happens. Science aims to provide understanding, as well as knowledge.
 - B. In contrast, empiricists have tended to think of science as constrained by and concerned with what happens. For some thinkers, the demand for explanation seems like an invitation to metaphysical speculation. Newtonians, for instance, felt no need to explain what gravity was; that seemed like a job for philosophers, not for scientists.
 - C. If one is not careful, explanations can collapse into verbal emptiness or expand into metaphysical excess.
 - D. For reasons such as these, some empiricists have taken the extreme-sounding measure of denying that science is in the explanation business. Scientific laws, such as Kepler's laws of planetary motion, are economical ways of describing experience. But it is no part of science to tell us why things happen.
- II. Carl Hempel's *covering-law model* of explanation is one of the great achievements of logical positivism. Hempel tries to reconcile empiricist scruples with the need for genuine scientific explanations.
 - A. Hempel links explanation and understanding by claiming that a complete explanation shows that the explained event or fact *had* to happen. We understand when we know that something *must* be the case.
 - B. But, as empiricists such as Hume have emphasized, experience provides no direct evidence of things “having to happen.” We experience no connections between events such that one makes the other happen.
 - C. Hempel solves this problem by appealing to logical necessity, the only notion of necessity that the logical positivists found clear and useful. For Hempel, explanations are arguments, and the truth of the premises necessitates the truth of the conclusion.
 - D. As is characteristic of logical positivism, Hempel offers a rational reconstruction of scientific explanation. He is not describing the explanations actually given by scientists; he is more interested in illuminating the logic of explanation than the practice of giving explanations.
- III. Testable laws of nature form the centerpiece of covering-law explanations (hence the name).
 - A. Hempel allows for two different kinds of *explananda*: laws and events.
 1. To explain a law is to derive it from other, more general laws. Thus, Newton's laws of motion explain Kepler's laws of planetary motion.
 2. More commonly, we explain events. To explain an event is to derive it from relevant laws combined with suitable initial conditions. Chemical and physical laws combined with facts about a match, the surface on which it was struck, the presence of oxygen, and so on explain its lighting.
 - B. The requirement that an explanation (non-trivially) contain testable empirical laws ensures that explanations will be scientific, rather than metaphysical.

1. Why is it acceptable to invoke magnetism to explain why iron behaves so differently from wood but not acceptable to invoke a life force to explain why living things behave so differently from nonliving things? The former explanation doesn't just posit an unobservable entity; it provides independently testable laws about the behavior of observable entities. The life force does not have any predictive content but is invoked only after the fact to explain things.
2. These independently testable laws link explanation and prediction very tightly for Hempel. Every adequate explanation is a potential prediction and every adequate prediction is a potential explanation. This symmetry between explanation and prediction links the covering-law model of explanation to the uncontroversial empiricist goal of prediction.
3. Explanations that meet the standards of the covering-law model provide the resources we need to both control and predict our experience. If I know a law such as that water expands when it freezes and I know how much water is in my radiator, then I know under what conditions my radiator will burst.

IV. Let's grant for now that the covering-law model's conditions are sufficient for a scientific explanation. Are these conditions necessary?

- A. There is some *prima facie* reason to think these conditions necessary. Less stringent conceptions of explanation (for example, reducing the unfamiliar to the familiar) face serious problems.
- B. One might, however, worry that Hempel's model rules out legitimate scientific explanations.
 1. Some have claimed that biological explanations at least sometimes proceed without appealing to laws of nature (for example, traits are explained by their functions, or events are explained by being situated in a narrative).
 2. In psychology or history, people's behavior is sometimes explained by reconstructing their goals or reasons. Arguably, such explanation involves no laws of nature.
- C. Hempel can offer one of several responses, depending on the circumstances of the example in question.
 1. Hempel has no problem admitting that some complete explanations are incompletely stated. You can explain why ice floats on water by saying that it expands when it freezes. Much of the explanation is unstated, but that's not usually a problem with the explanation. But the cases from biology, psychology, or history arguably wouldn't include laws even if the explanations were stated completely.
 2. Hempel also makes room for a notion of partial explanation. Insofar as evolutionary biology allows for a prediction that a species of a certain description will emerge in given circumstances, it can explain the existence of a species of that type. Perhaps it explains the existence of a small scavenger, for example, but not of a weasel. But this still imposes major restrictions on the explanatory aspirations and power of biology.
 3. Hempel can allow that a narrative provides resources for explanation, but not that, by itself, it can constitute an explanation. Thus, the story of evolution, as opposed to the theory of evolution, provides no explanation at all.
 4. Even the theory of evolution, Hempel must insist, explains relatively little. What a theory would not have been in a position to predict, it is not in a position to explain. And biological phenomena involve so much complexity and randomness that biology can offer only vague and probabilistic predictions or explanations.
 5. Hempel handles psychology and history similarly. At best, given the state of laws in these fields, we can muster partial and probabilistic explanations.
 6. As impressive as Hempel's model is, one must ask whether we should be willing to pay the price it demands by excluding so many explanations from biology and other sciences.

Essential Reading:

Hempel, "Laws and Their Role in Scientific Explanation," in Boyd, Gasper, and Trout, *The Philosophy of Science*, pp. 299–315.

Supplementary Reading:

Rosenberg, *Philosophy of Science: A Contemporary Introduction*, chapter 2.

Questions to Consider:

1. Anyone who has spent time with a toddler knows that one can ask “why” about a great many things. When is explanation called for? Should science (or philosophy) be in the business of explaining everything (for example, are we supposed to explain why there is something rather than nothing)? If not, how are we to decide which “why” questions are badly posed?
2. Do you think that we offer reasonable approximations to scientific explanations when we explain each other’s behavior, or do you think we fall short of that standard? If we fall short, does the problem lie with our explanations or with the standard or both?

Lecture Nineteen

Putting the Cause Back in “Because”

Scope: Though it ruled the explanatory roost for quite some time, the covering-law model faces very serious problems. It seems committed to allowing that Mr. Jones, having taken his wife’s birth-control pills, explains his failure to get pregnant. That, to put it mildly, seems unfortunate. The *causal-relevance conception* of explanation is now preeminent, but it, too, faces challenges, notably the fact that causation is a notoriously tricky concept about which to get clear. In addition, some explanations are dubiously causal and some seem clearly non-causal. How much better has our theory of explanation gotten?

Outline

- I. We saw some reasons last time to worry that Hempel’s covering-law model of explanation might be too restrictive. But the more serious worry is that it is too permissive. It counts arguments that intuitively have no explanatory force as legitimate scientific explanations.
 - A. The covering-law model allows explanations of causes by effects or by symptoms.
 1. The same laws that allow us to infer the length of a flagpole’s shadow from the height of the flagpole also allow us to deduce the height of the pole from the length of the shadow.
 2. But we tend to think that explanation is an asymmetric relation: We think that the height of the flagpole explains the length of the shadow and that the length of the shadow does not explain the height of the flagpole.
 3. For similar reasons, Hempel’s model allows symptoms to explain the things for which they are symptoms. It seems right to say that the barometer is falling because a storm is approaching. But are we comfortable saying that a storm is approaching because the barometer is falling?
 - B. Hempel’s model also permits intuitively “wrong-way” explanations with respect to time. We can explain a planet’s future location in the sky by appealing to its present location and some laws of planetary motion. But can we explain its present location by appealing to its future location plus the same laws?
 - C. Further, Hempel’s model seems to allow for irrelevant explanations. If Mr. Jones takes his wife’s birth-control pills, we can certainly predict, using laws of nature, that he will not become pregnant. But have we explained this fact?
- II. Many philosophers appeal to causation to avoid problems like those just noted. The *causal model* of explanation, simply stated, says that to explain an event or fact is to provide information about its causes.
 - A. The covering-law model gets into trouble because the notion of expectability on which it relies is too symmetrical. Causation provides a needed asymmetry. The height of the flagpole produces the length of the shadow, and the approach of the storm causes the falling barometer reading. The past causes the future but not vice versa. Explanation tracks causation.
 - B. Explanations that include irrelevant information fail because they lead away from the actual causes. It’s the fact that Mr. Jones is male, rather than the fact that he takes birth-control pills, that causes (and, hence, explains) his failure to become pregnant.
 - C. The covering-law theorist has some resources for accommodating causal intuitions within the covering-law model, but most philosophers see serious problems here.
 - D. Once we move to the causal model, we arguably do not need the whole covering-law apparatus or arguments, laws, and so on. An event can be explained simply by saying what caused it.
- III. The biggest problem facing the causal model involves figuring out just what causation amounts to.
 - A. Empiricists, such as Hume and Hempel, are suspicious of causation. They note that we observe correlations but not causation and insist that causal talk get cashed out in experiential terms. We will keep these empiricist scruples in mind as we examine some influential accounts of causation.

- B. Many find it natural to think of causation as involving a kind of physical connection, a transfer of something (for example, momentum) from the cause to the effect. Though this conception of causation is intuitive, it has many counterintuitive consequences.
 - 1. It has problems counting absences as causes. It might not count drowning as a cause of death, because it is the absence of oxygen that causes death.
 - 2. Some potential cases of causation do not seem to be connected in space and time the way this approach requires. If we think the death of Socrates causes Xantippe to become a widow, then we have to allow that causation travels, as it were, instantaneously across space.
 - 3. Conversely, it seems counterintuitive to allow just any relevant absence or omission to count as a cause. Did my failure to throw a rock at a window cause the window not to break?
 - C. *Regularity* theories of causation are popular among those who have empiricist scruples.
 - 1. The most common such view says that a cause is a necessary part of a condition that, together with the laws of nature, is sufficient (but not necessary—we are looking for *a* cause, not *the* cause) for its effect. Thus, the presence of oxygen counts as a cause of the match lighting in much the same way that my striking the match counts.
 - 2. Sometimes, we pick out one cause as special and talk as if it is *the* cause. If you leave your iron on and your house catches fire, we say that the iron, not the presence of oxygen in the atmosphere, caused the fire, but that's not strictly true.
 - D. *Counterfactual* approaches analyze causation in terms of what would have happened had things gone otherwise. Because the match would not have lit had I not struck it, the striking is a cause of the lighting.
 - E. The regularity view might be too empiricist, and the counterfactual view might not be empiricist enough.
 - 1. The regularity approach is empiricist-friendly, because the only connection between cause and effect is logical, not physical. But for this reason, it runs into problems like those plaguing the covering-law model. It looks as if, given the laws of nature, falling barometers cause storms, because we can derive storms from falling barometers and laws, but we cannot perform the derivation if the falling barometer is not included (thus, the barometer is a necessary part of a sufficient condition).
 - 2. We've noted before that empiricists are uncomfortable with counterfactuals. They don't like talk of how things would have been if the barometer hadn't fallen.
 - F. Cases of overdetermination make mischief for most accounts of causation. Suppose that two sharpshooters fire at the same time and accurately at a condemned prisoner.
 - 1. On a standard regularity analysis, both shooters cause the prisoner's death. Each shot is a necessary part of a sufficient condition of the prisoner's death.
 - 2. According to a simple counterfactual view, neither shooter caused death, because it is true of each of them that, had he not pulled the trigger, the prisoner would have died anyway.
 - 3. Both the regularity and counterfactual approaches capture some of our intuitions about causation, and neither captures all such intuitions. This is because our notion of causation is, at best, tricky and complicated.
 - G. Cases involving preempting causes can vex both of these views as well. Suppose Jones eats a pound of arsenic but then gets run over by a bus before the arsenic takes effect.
 - 1. As always, there's room for more sophistication than we can do justice to, but a simple regularity view will count both the arsenic and the bus as causes of Jones's death.
 - 2. And a simple counterfactual view will say that neither event caused Jones's death.
 - H. Finally, it is worth noting that causation is not transitive. X can cause Y, which causes Z, without it being the case that X caused Z.
- IV. Returning to explanation and waiving these problems about the notion of causation, the causal model has to face the challenge that other views have encountered, namely, does it include illegitimate scientific explanations or exclude legitimate scientific explanations?
- A. Many standard views of causation have the consequence that the complete causal history of an event comprises its full cause and, hence, according to the causal model, its explanation.

1. The model thus looks too permissive if it allows the Big Bang to count as a cause (and, hence, as an explanation) of the fact that I'm giving this lecture.
 2. Advocates of the causal model can distinguish between the true and the useful here. It is true, strictly speaking, that the Big Bang explains my giving this lecture, but it's not a helpful explanation to give, which is why it sounds absurd to us.
- B.** Laws do not cause other laws to be true; thus, the causal model will need supplementation if it is to handle explanation of laws. This problem cannot be handled by the causal approach, but it is all right, perhaps, to have different accounts of explanation for laws and for events.
- C.** Some explanations seem to proceed by identification, and that looks incompatible with causation. It appears that the average kinetic energy of the molecules of a gas sample can explain its temperature.
1. Arguably, this is a case of a fact explaining itself. We'll discuss such cases in an upcoming lecture.
 2. But because no fact can cause itself, the explanation is non-causal.
- D.** We saw that the covering-law model seemed to give short shrift to biological explanations. Is the causal model any friendlier to biological explanations?
1. Let's focus on one important subclass of biological explanation—*functional* explanation. Why do mammals have hearts? "For pumping blood" seems like a decent explanation.
 2. The covering-law model has a problem here, because there is no law saying, for example, that whenever a species needs blood pumped, it will develop a heart.
 3. The causal model would seem to face a problem here as well, because to describe what something is for or what it does seems very different from describing how it was brought about. But important work has been done in recent decades to show how an explanation such as "mammals have hearts for pumping blood" can be construed as a causal explanation in terms of evolutionary history. The idea is that the existence of a given heart in a given mammal is explained by the causal contribution to the reproductive fitness of the creature's ancestors that past hearts have made.

Essential Reading:

Ruben, "Arguments, Laws and Explanation," in Curd and Cover, *Philosophy of Science: The Central Issues*, pp. 720–745.

Supplementary Reading:

Mackie, "Causes and Conditions," in Brody and Grandy, *Readings in the Philosophy of Science*, pp. 235–247.

Questions to Consider:

1. Does causation need to be some kind of physical process or of a certain magnitude to be scientifically legitimate?
2. A few philosophers have held that we sometimes directly observe causation. How plausible do you find such a claim?

Lecture Twenty

Probability, Pragmatics, and Unification

Scope: In this lecture, we examine the remaining major issues in the philosophy of explanation. We sketch the main competitor to causal accounts, according to which explanation is achieved by “doing the most with the least” by unifying diverse phenomena under a small number of patterns and principles. We then consider the radical proposal that explanation is no part of science itself and that good explanations are nothing deeper than contextually appropriate answers to “why” questions. Finally, we examine the major accounts of statistical explanation and ask whether there can be explanations for irreducibly probabilistic phenomena.

Outline

- I. The leading idea behind *unificationist models* of explanation is that scientific explanation increases our understanding by reducing the number of independent explainers we need. The fewer primitive principles and styles of argument we need to posit, the more unified and the more explanatory is our science.
 - A. The central challenge here is to figure out what unification amounts to.
 1. It is not enough for a theory to imply a bunch of statements. “Ice floats in water; copper conducts electricity; and bears are mammals” implies each of the smaller statements of which it is composed, but it achieves no unification.
 2. A more promising idea says that a theory unifies when it minimizes the number of statements that are treated as independently acceptable. Newton’s physics allows us, at the cost of adding a few independently accepted law statements, to start from a modest number of initial conditions and derive, rather than posit, countless other statements about how things move. In this sense, Newton unifies by helping us do the most with the least.
 3. A somewhat similar approach tries to minimize argument patterns. The reason that birth-control pills do not figure in an explanation of Mr. Jones’s failure to get pregnant is that we have a simpler, more unified theory if we appeal to arguments involving males not getting pregnant than we do if we appeal to arguments involving birth-control-pill-taking males not getting pregnant.
 4. Like the covering-law model, this approach tries to get logical relationships to do the work done by metaphysical relationships in the causal model. The idea is that we *systematize* our arguments in such a way that we can get the most out of them, and an argument counts as an explanation if it figures in the best systematization of our theories.
 - B. Like its competitors, the unification model faces significant challenges.
 1. Relatively local unification, such as breaking a code, hooks up very nicely with understanding and explanation. It’s less obvious that global unification bears the same relationship to understanding and explanation.
 2. Some philosophers claim it is possible to unify causes in terms of effects rather than effects in terms of causes, and this brings us back to some of the counterintuitive features of the covering-law model. Will any sort of logical relation capture some of the asymmetries that seem essential to explanation?
- II. Bas van Fraassen denies that there is a correct account of scientific explanation as such. For him, an explanation is merely an answer to a “why” question.
 - A. Which question is being asked and what counts as a good answer to it depend on context.
 1. “Why” questions typically assume an implicit contrast. The bank robber Willy Sutton’s priest meant to ask him, “Why do you rob banks rather than have a job?” but Sutton took “Why do you rob banks?” to mean “Why do you rob banks rather than other places?” He replied: “Because that’s where the money is.” Sutton did not give a good explanation because he did not give a good answer to his interlocutor’s question.
 2. Good answers will take the interests, abilities, and information of the audience into account. A perfectly correct quantum mechanical explanation of why a square peg won’t fit in a round hole is still a bad explanation if offered to a 5-year-old.

3. And there is no noncontextual standard of explanatory goodness; causal or unifying or covering-law explanations can all be good ones in the right context, and so can many other kinds of explanations.
 4. Van Fraassen thus repudiates the ambitions of such thinkers as Hempel, for whom explanation is no more contextual than mathematical proof is; a good proof given to a 5-year-old is still a good proof.
- B. For van Fraassen, not only is there no distinctively scientific notion of or standard for explanation, explanation is itself no part of science.
1. We use science in giving explanations, but we're not doing science when we explain. Explanation is rather like technology in this respect.
 2. Van Fraassen's reasons for this are empiricist ones; if the demand for explanation is built into science, it will lead inexorably to metaphysics.
- C. Like all the other views we've examined, van Fraassen's position faces significant criticisms.
1. It is difficult to specify what makes a given answer relevant to a given question. If one isn't careful, any answer counts as relevant to any question.
 2. Do we really want to give up on the idea that science has distinctive explanatory goals and standards of explanatory adequacy?

III. *Statistical explanation* is of independent importance and raises its own distinctive set of problems.

- A. We might resort to statistical explanation in either of two very different circumstances.
1. Because we might not have enough information about a situation to explain it deterministically, we settle for the statistical claim.
 2. Alternatively, the situation might be irreducibly indeterministic. The dominant interpretation of quantum mechanics says that it is just a brute fact about the universe that a uranium-238 atom has a certain probability of decaying in a given time period.
- B. The covering-law model provides the classic account of statistical explanation.
1. In Hempel's account, statistical explanations use statistical laws and initial conditions to confer a high probability on the explanandum.
 2. The statistical law refers to an objective probability, some kind of fact in the world.
 3. The probability that is conferred on the explanandum is something different. It is either a logical statement about the amount of evidential support premises confer on a conclusion, or it is a personal probability, an estimate or degree of belief about this evidential support.
 4. Despite clear similarities to the deterministic covering-law model, the statistical case is very different. Because the argument is inductive, additional information can make a difference to the probability conferred on the explanandum. If we know that Jones has a particular kind of infection and has been given penicillin, we might be able to cite a statistical law saying that a high percentage of people with this infection who get penicillin recover within 24 hours. But if we add the information that Jones is allergic to penicillin, the relevant statistical laws will change considerably.
 5. Thus, on the covering-law model, explanation is relative to an information situation. This relativization to an information situation makes the notion of a good statistical explanation problematic. Do we want to say that we have a good statistical explanation only when all the relevant information is included in the explanation? Or do we instead want to say that we have a good statistical explanation if there is no *known* additional information that would change probabilities? Or do we want to count as good any explanation that uses true statistical laws to derive a probability for Jones's recovery?
- C. Surprisingly, the probability conferred on the explanandum might not have much to do with whether or how well it has been explained.
1. If I know that I have a red-biased roulette wheel, don't I understand why it sometimes comes up black just as well as I understand why it generally comes up red?
 2. This suggests that it is not necessary for an explanation to confer a high probability on the explanandum. If a person who has a disease that is invariably fatal if untreated undergoes a procedure that has a 30% chance of curing the disease, then the procedure explains the person's survival even though it doesn't make survival particularly likely. Thus, it's not necessary that an explanation confer high probability on the explanandum.

3. Nor is it sufficient for a good explanation that the *explanans* confer a high probability on the explanandum. Problems of irrelevance, analogous to those in the deterministic case, arise. If you take vitamin C, you will recover from a cold within seven days. But given that you will recover within seven days even if you do not take vitamin C, the explanation is not a good one.
- D. The covering-law model makes explanation a matter of an argument that renders the explanandum highly probable. A competitor, analogous to the causal model, suggests that explanation is a matter of raised probability, rather than high probability. It is even possible that a cause might be the kind of event that lowers the probability of the outcome happening. But for a causal theorist, causes explain, even when they render the explanandum less likely than it had been. Thus, explanation would amount to probabilistic relevance rather than probability-raising.

Essential Reading:

Hempel, "Inductive-Statistical Explanation," in Curd and Cover, *Philosophy of Science: The Central Issues*, pp. 706–719.

Van Fraassen, "The Pragmatics of Explanation," in Boyd, Gasper, and Trout, *The Philosophy of Science*, pp. 317–327 (also in Balashov and Rosenberg, *Philosophy of Science: Contemporary Readings*, pp. 56–70).

Supplementary Reading:

Rosenberg, *Philosophy of Science: A Contemporary Introduction*, chapter 3.

Godfrey-Smith, *Theory and Reality: An Introduction to the Philosophy of Science*, chapter 13.

Questions to Consider:

1. Disputes such as the one we've seen in this lecture between contextualists, such as van Fraassen, and their critics crop up everywhere in philosophy. Generally speaking, contextualists think that noncontextualists are searching for grander answers than the phenomena will permit, while noncontextualists think that contextualists give up on the proper ambitions of theorizing too easily. Which side of this controversy gets your sympathy and why?
2. Does it seem more reasonable to you to say that objectively improbable events are inexplicable or to say that we can explain them, even though they are objectively improbable?

Lecture Twenty-One

Laws and Regularities

Scope: Empiricists such as Hempel rely on laws of nature in their accounts of explanation. But the notion of a natural law has occasioned a great deal of suspicion from empiricists. Suppose the statement “All the beer in my refrigerator is American-made” is true. Nevertheless, we think that it could easily have been false. Contrast this with a statement such as “All copper conducts electricity.” How, without appealing to claims that are not properly funded by experience, can one explain why the second of these true statements expresses a law and the first does not? Empiricists have tackled this problem very resourcefully, but we know by now that it won’t be easy for them.

Outline

- I. We have seen that laws of nature figure centrally in Hempel’s approach to explanation. Quite apart from the merits of Hempel’s approach, laws of nature are very much worthy of philosophical attention in their own right.
 - A. It is generally, though by no means unanimously, agreed that science seeks to uncover laws of nature. The role of such laws in various sciences is a matter of considerable controversy.
 - B. The notion of a law of nature, like those of explanation and causation, has seemed suspicious to empiricist philosophers. It has had associations with divine decrees and other metaphysical pictures.
 - C. Laws of nature are to be distinguished from positive laws (what lawyers study) and from logical laws (which are analytically and necessarily true). Laws of nature are synthetic and contingent. God could have set up the universe with gravity inversely proportional to the cube of distance instead of the square. This idea of a contingent rule that objects and events are somehow bound to follow bothered Ayer and other empiricists.
 - D. Most laws of nature are of universal conditional form: “All As are Bs.” Many laws (such as the law of supply and demand) that do not appear to have this logical form actually do.
 - E. It is generally (but not unanimously) agreed that a statement cannot be a law of nature unless it is true.
 - F. Some statements called laws are not true and, thus, are not really laws (for example, Newton’s gravitational law gets corrected by general relativity). Many laws are called *equations*.
- II. *Regularity accounts* of laws of nature treat them as statements about what always happens. They are patterns in experience rather than something above and beyond patterns that govern or control events. These are empiricist-friendly theories, but simple versions of them face devastating problems.
 - A. The simplest version of a regularity account says that any true, contingent statement of universal conditional form is a law of nature.
 - B. Such an approach cannot distinguish laws from *accidental generalizations*. Although it may be true that “All the beer in my refrigerator is American-made,” the statement does not seem to express a law of nature, even though it has the right logical form.
 1. Even if it always has been and always will be true, the statement still doesn’t seem like a law of nature.
 2. We cannot disqualify the statement on the basis of its being restricted to just one place (that is, my refrigerator). We want to allow for the possibility that some laws hold at just one place (such as Earth).
 - C. *Vacuous laws*—laws that do not have instances—also present a problem for regularity accounts.
 1. “All particles that travel faster than the speed of light are pink” is logically equivalent to “No particle traveling faster than the speed of light is non-pink.” Once you realize “All dragons like jazz” is equivalent to “There are no jazz-hating dragons,” it is not so hard to admit the former statement to be true.
 2. Because the statement about pink particles is true (and fully general), it looks as if it should count as a law.

3. The simplest way to fix this would be to require that there be at least one particle traveling faster than the speed of light. But not all laws that lack instances are illegitimate. The fact that there are no bodies on which no forces act does not prevent Newton's first law from being a law.

III. *Epistemic regularity accounts* distinguish laws from other generalizations on the basis of how we treat them. This approach can handle the problems noted above, but it faces a major problem of its own.

- A. We think that laws, but not accidental generalizations, tend to support counterfactuals. If my pen were made of copper, it would conduct electricity, but we do not think that if a German beer were placed in my refrigerator, it would become an American beer.
- B. Laws, as we saw with Hempel, hold a special place in our explanatory practices. We explain why an object conducts electricity by saying that it is made of copper, but we do not explain why a given beer is American by saying that it came from my refrigerator.
- C. Laws are relatively central to our webs of belief. They are not easily undermined by new information or falsified by putative counterexamples. If I'm told that this is a freshly purified piece of copper, I will still think it conducts electricity. If I'm told that there's a new brand of beer in my refrigerator, I might well doubt that it is American.
- D. Laws are more readily confirmed by their instances than are accidental generalizations. After a modest number of samples, you are convinced that copper conducts electricity. But even if you know a number of weird philosophers, you might resist the generalization that all philosophers are weird. You are more confident that you might find a non-weird philosopher than that you might find some nonconductive copper.
- E. The problem with this analysis is that it does not make room for undiscovered laws. On this view, the laws of nature are what they are because of how we handle them.

IV. The *systems theory* is the most sophisticated of the broadly empiricist approaches to laws of nature.

- A. The laws of nature flow from deep structural patterns in actual events. We identify the patterns by looking at the best ways of describing the world in a deductive system. *Best* means simplest and strongest.
 1. Simplicity is inversely proportional to the number and complexity of the axioms. "Everything happens according to the will of Elvis" is a simple theorem, with one object and one overall argument scheme. You might need more independently acceptable sentences in order to flesh out *will*, but it's still a pretty simple system. At the other extreme, treating every event that happens as an axiom, as an independently acceptable sentence, makes for a highly "unsimple" system.
 2. Strength is a matter of how informative the theorems are. The list of all the phenomena is strong, but the Elvis theory is not. The list of axioms tells us (albeit inconveniently) what to expect, while the Elvis theory is convenient but gives us little to go on.
 3. Simplicity and strength thus work against one another. It is easy to have a simple theory that does not say much or a strong theory that is not simple.
- B. Laws of nature, according to systems theorists, are all the true, contingent generalizations that figure in all the best deductive systems.
 1. "There are no uranium spheres a mile or more in diameter" is plausibly a law. Our best physical theories will belong to the best deductive systems, and they imply that a quantity of uranium like that would explode.
 2. "There are no gold spheres a mile or more in diameter" is not a law of nature, even on the assumption that it is true. It does not follow from our best physical theories. We could add it as an axiom, but that would lessen the system's simplicity without a compensatory payoff in strength.
- C. This looks like a promising way of handling the problems that plagued the simple regularity view and the epistemic regularity view.
 1. Laws that do not have instances are allowed, but they have to pay their way in the currency of simplicity and strength.

2. Laws can be restricted in space or time but, again, only if they pay off in terms of simplicity and strength. There might be laws that apply only to Earth, but there won't be laws that apply only to my refrigerator.
3. Arguably, this view can explain why the laws of nature figure as they do in explanation, why they support counterfactual conditionals, and so on.
4. On the systems approach, laws are not constituted by our handling them a certain way. Thus, the systems approach can make room for undiscovered laws.

Essential Reading:

Ayer, "What Is a Law of Nature?" in Curd and Cover, *Philosophy of Science: The Central Issues*, pp. 808–825.

Supplementary Reading:

Earman, "Laws of Nature," in Balashov and Rosenberg, *Philosophy of Science: Contemporary Readings*, pp. 115–126.

Questions to Consider:

1. Does it seem sensible or unfortunate to you that we use the term *law* to refer both to positive laws and to laws of nature (not to mention logical and mathematical laws)? What similarities are highlighted by this term and what differences are obscured by it?
2. Some crucial laws seem not to have any actual instances (for example, a law describing how objects move on a frictionless plane). How do we come to know such laws?

Lecture Twenty-Two

Laws and Necessity

Scope: Philosophers who are not burdened by the empiricist scruples that motivate regularity theorists are free to develop more “metaphysical” conceptions of laws. But, of course, these bring with them their own challenges. According to such views, laws describe powers, dispositions, tendencies, or relations among properties. These conceptions put us in a position to explain why a given piece of copper *must* conduct electricity. It is not clear, however, how we’re in a position to gain knowledge about such tendencies and relations among properties. Furthermore, these views might require us to accept the somewhat uncomfortable doctrine that there are facts in the world that aren’t fixed by all the particular facts of the world. In other words, do we want to allow for the possibility that two identical worlds could be governed by two different sets of laws?

Outline

- I. For philosophers who reject the constraints of a Hume-inspired empiricism, laws of nature look very different than they do to those who respect such constraints.
 - A. For empiricists, laws do not say anything different than accidental generalizations do. A law is a true generalization that is used in a certain way or that fits into a certain system.
 - B. *Necessitarians*, on the other hand, take laws to describe tendencies or powers, not themselves directly observable, that explain observable phenomena.
 1. Empiricists tend to see laws as asserting relationships that hold among objects: All objects that are made of copper are objects that conduct.
 2. Necessitarians tend to see laws as having a different logical form. A law does not assert a relationship between objects. The version we will consider says that a law asserts a relationship between properties—being made of copper *makes for* conducting electricity. For that reason, if something is made of copper, it *must* conduct electricity.
 - C. Necessitarians grant that experience shows, at best, that all copper conducts electricity, not that it must. But they think that empiricist scruples must be set aside if there are to be statements worth calling laws of nature. If all we have are generalizations about objects, then we have no laws, because the generalization tells you only that you won’t find nonconductive copper, not that you can’t find it. Laws do not just describe; they govern.
 - D. For the positivists and Quine, all necessity is linguistic necessity. Necessitarians hold that laws assert relationships of physical necessity.
 - E. The law itself, however, is contingent. Constitutional law provides a useful example here. Given the contingent fact that our Constitution says that the president must consult Congress before declaring war, the president *must* do so. Similarly, “copperness” need not have necessitated electrical conductivity (the universe could have been built differently), but because it does so necessitate, a given piece of copper must conduct electricity.
- II. Like the systems approach from our last lecture, the necessitarian conception can handle the problems that plagued the simple and epistemic regularity accounts of laws.
 - A. Vacuous laws are no problem. There is no relation of necessity between going faster than the speed of light and being pink; thus, “All particles that exceed the speed of light are pink” does not express a law. But there can be legitimate laws without instances. In a Newtonian universe, there is a relationship of necessity between the (uninstantiated) property of being an object on which no forces act and the property of having zero acceleration.
 - B. The difference between laws and accidental generalizations is quite stark on such a view. Being made of uranium necessitates being less than a mile in diameter; being made of gold does not.
 - C. It is clear that, for a necessitarian, laws are discovered, not made.
 - D. These relations of necessity explain why laws support counterfactuals.

1. It is hard to see how we get from “All copper conducts electricity” to “If this were made of copper, it would conduct electricity.” The former is a statement about what always happens in this world; the latter is a statement about what would happen in a different possible world.
 2. But it is easy to see how we get from “Being made of copper necessitates conducting electricity” to “If this were made of copper, it would conduct electricity.” Relations among properties carry over across possible worlds relatively straightforwardly.
- E.** This same necessity explains explanation. Necessitarians deny that even the systems approach provides a good account of explanation.
1. The fact that all copper conducts electricity implies that this piece of copper conducts electricity. But to imply is not to explain. And, necessitarians add, to be implied by a strong and simple deductive system is not yet to be explained.
 2. Laws can explain precisely because they are not universal generalizations. The relation between properties explains why certain relations hold among objects.
- III.** Empiricists, unsurprisingly, find the necessitarian approach objectionable. Empiricists and necessitarians disagree about the epistemology and about the metaphysics of this necessity relation.
- A.** Empiricists want to know just what this relation of making-necessary is.
1. How do we tell when we have a case of this necessity relation? How do happenings reveal what must happen?
 2. What grounds this necessity relation? We have a rough idea of what grounds the “must” in “The president must consult Congress,” but do we have an analogous basis for “Copper must conduct electricity”?
- B.** Empiricists and necessitarians disagree about whether laws can “float free” of the particular facts about what happens.
1. The empiricist maintains that we have to build our laws and our conception of what is physically possible out of what is physically actual. If two worlds had the same facts but different laws, we would have no factual basis for determining the laws. Laws would outrun any factual constraint on them.
 2. The necessitarian replies with arguments designed to show that the facts about particulars and particular happenings are insufficient to determine the laws. Because we cannot do without laws, we need to drop the idea that all the facts are determined by particular facts. Necessitarians think there could be laws governing kinds of things that never happen.
- IV.** Nancy Cartwright, a philosopher of physics, argues for a stark dilemma. Either the laws of nature are false but can be used in scientific explanations, or they are true but useless for explaining things.
- A.** Our most fundamental laws do not describe how bodies actually behave. The law of gravitational attraction says only how bodies would move if only gravitational forces were acting on them. But that is hardly ever true because of electrical charges and so on.
- B.** For the law to be true, it would either have to be tremendously restricted (that is, to those bodies on which no other forces are acting) or “hedged”—protected by a powerful “all-other-things-equal” clause. But either of these strategies threatens to rob the law of explanatory power—we need the law to apply in some sense to bodies whose motion it is supposed to explain.
- C.** Roughly speaking, Cartwright thinks that her argument supports a necessitarian construal of laws: They describe not what actually happens but tendencies or powers that explain what happens. What laws are true “of” is not the observable behavior of objects.
- D.** Cartwright’s critics have offered some interesting competing proposals.
1. They have construed laws as describing actual component forces, rather than potentialities, powers, or tendencies.
 2. They have argued that context allows us to preserve the truth of unhedged laws. If I say that this house is empty, you do not falsify my claim by pointing out that there are light bulbs in the fixtures. And you don’t falsify a law such as “Metal bars expand when heated” by pointing out that they don’t if someone is hammering on both ends of them.

Essential Reading:

Dretske, "Laws of Nature," in Curd and Cover, *Philosophy of Science: The Central Issues*, pp. 826–845.

Cartwright, "Do the Laws of Physics State the Facts?" in Curd and Cover, *Philosophy of Science: The Central Issues*, pp. 865–877.

Supplementary Reading:

Carroll, *Readings on Laws of Nature*.

Questions to Consider:

1. How intelligible do you find the idea of physical necessity? Does physical necessity require a kind of grounding like that of (positive) legal necessity? If so, what is it? If not, why not?
2. How, if at all, are "hedged" or "all-other-things-equal" laws empirically testable? Astrology is full of hedged "laws." Are similar laws in history, psychology, economics, or physics any better?

Lecture Twenty-Three

Reduction and Progress

Scope: Prominent episodes in the history of science seem to involve one theory being “absorbed by” or reduced to another. Kepler’s laws of planetary motion appear to reduce to Newtonian dynamics, and genetics seems to reduce to molecular biology. Similar points can be made in terms of entities (water reduces to H_2O). If one theory assimilates another and explains its phenomena in more fundamental terms, it seems that our understanding deepens and improves, and this kind of theory change seems straightforwardly progressive. We begin by examining the account of reduction given by the positivists, according to which *bridge principles* allow the reduced theory to be derived from the reducing theory. Kuhn and Feyerabend held that many cases identified as reductions by the positivists are more like *replacements* of one theory by another. To what extent does science progress through reductions, and how smooth is such progress as it makes?

Outline

- I. We turn to the issue of reduction, which will draw on some ideas we have been discussing, especially about explanation and scientific progress. A *reduction* takes place, speaking loosely, when something is shown to be “nothing but” something else.
 - A. Reduction can be applied to stuff (water reduces to H_2O), laws or theories (Kepler’s laws reduce to Newton’s), or to whole disciplines (some think that biology reduces to physics).
 - B. Many of us initially take such talk to be primarily concerned with ontology, with what exists. The positivists tended to be suspicious of this way of speaking. It seemed to them to invite metaphysical questions about how many things or kinds of things there are in the universe.
 - C. Accordingly, the positivists construed reduction linguistically. It is talk of water that reduces to talk of H_2O . The general question is: Under what conditions do certain theories, statements, or terms reduce to others? We focus on the case of theories.
 - D. If Quine’s holism (discussed in Lecture Eight) is right, this distinction between deciding *how to talk* and deciding *what there is* won’t hold water.
- II. We’ve discussed a couple of philosophical reductions in this course; scientific reductions are a bit different.
 - A. We’ve seen Berkeley try to reduce talk of objects to talk of ideas and experiences.
 - B. We’ve seen the positivists try to reduce talk of unobservables to talk of observables.
 - C. In philosophical reductions such as these, issues about what is knowable on the basis of *what* loom large and issues about meaning loom very large. Berkeley and the positivists think that language cannot meaningfully be deployed beyond the bounds of experience.
 - D. In science, one does find reductions driven by issues of meaning (for example, Bridgman’s operationalism), but philosophical concerns loom large in such reductions. Standard scientific reduction of one theory to another involves issues of explanation and progress more directly than issues of meaning.
- III. Reductions seem to constitute clear cases of scientific progress.
 - A. Reductions involve a kind of theory change that isn’t mere change. The old theory is not discarded as false; it is, instead, preserved within a richer theory.
 - B. Progress manifests itself in such things as increasing explanatory power. The reducing theory would now be able to *explain* facts and laws of the reduced theory. If classical thermodynamics reduces to statistical mechanics, then facts about temperature, for instance, are a special case of facts about molecular motion.
- IV. The classical positivist conception of reduction treats it as a deductive relationship. One theory reduces to another if the former can be derived from the latter. The reduced theory is, as it were, logically contained in the reducing theory.
 - A. If the reduction is *homogeneous*, the case is relatively straightforward and relatively uninteresting.
 1. A reduction is homogeneous when the terms of the reduced theory are present in or can be defined using standard logical operations on the terms of the reducing theory.

2. The reduction of Galileo's and Kepler's laws of motion to Newton's more general laws of motion provides the classic case. Such terms as *velocity* appear (with the same meaning) in all these theories; Newton just unifies terrestrial and celestial motion.
- B. The more interesting reductions are *heterogeneous* and require *bridge principles* (also known as *bridge laws*).
1. Typically, the reduced theory uses terms that are foreign to the reducing theory. *Heat* and *temperature* do not figure in statistical mechanics.
 2. For this reason, it is not clear how thermodynamics can be derived from statistical mechanics. Statements about molecules in motion won't, without conceptual enrichment, get you statements about heat and temperature.
 3. Thus, principles connecting the terms of the reduced theory and the reducing theory must be added to the reducing theory. Such statements (for example, "Temperature is mean molecular kinetic energy") are *bridge principles* or *bridge laws*. The reduced theory is supposed to follow from the reducing theory plus bridge principles, not from the reducing theory alone.
- C. What kinds of claims are bridge principles?
1. They are not plausibly considered definitions: *Temperature* does not *mean* the same thing as *mean molecular kinetic energy*. For this reason, the reduction, though characterized in terms of logical/linguistic relationships between reducing and reduced theory, is not a semantic reduction. We are not reducing the content of one theory to that of the other.
 2. Bridge principles, then, probably should be considered empirical hypotheses that *identify* objects or processes. Being a donor of a single electron (a property in the language of physics) is both necessary and sufficient for having a valence of +1 (a property in the language of chemistry). The necessary and sufficient conditions could come apart, but that won't concern us.
 3. For purposes of reduction, many philosophers think that bridge principles that are much weaker than identifications will suffice; for example, such a weak principle might encompass a sufficient condition for the reduced property in the language of the reducing property. This is plausible, but we will focus on the simpler case of identification, which might be needed for some of the grander purposes of reduction.
- V. This classical positivist approach to reduction faces several challenges.
- A. Because nothing inconsistent with a theory can be derived from it, the statements of the reduced and the reducing theories must be logically consistent.
- B. But even in the least problematic cases of reduction, the reducing theory generally *corrects* the reduced theory. Galileo's law attributes constant acceleration to falling bodies near the Earth's surface, while Newton has the acceleration vary with the body's distance from the Earth's center of mass. Because Galileo's theory is incompatible with Newton's, the former cannot be derived from the latter.
- C. The classic response is that an *approximation* of Galileo's law can be derived from Newton's theory.
1. But we should admit that we have not reduced the theory, only a suitably corrected version of it.
 2. We now must face the question of how different reducing a theory is from replacing it.
 3. In addition, it turns out to be difficult to give a clear sense to the notion of an approximation that covers the range of reductions that take place in science.
- D. The holism of Feyerabend and Kuhn would have it that the meanings of key terms rarely remain constant across major theory change.
1. This point, even if granted, need not be fatal to the reductive project. Property identifications will allow some derivations to go through. Even though *temperature* doesn't have same meaning as *mean molecular kinetic energy*, we can substitute the one expression for the other in scientific laws.
 2. But Kuhn and Feyerabend think that the incommensurability of theories generally prevents such identifications.
- E. Even if one does not adopt so radical a move, however, clear problems of meaning and reference arise.
1. It is often thought, for example, that genes reduce to sequences of DNA, thus allowing classical genetics to reduce to molecular genetics. But the classical notion of a gene was characterized in three different ways, and each way corresponds to a different DNA segment. To what, if anything, has the classical notion of a gene been reduced?

2. It is tempting to appeal to a notion of approximation here, but it is difficult to find a generally acceptable account of approximation.
- F. We need enough flexibility to allow for corrective reductions, but we do not want our account of reduction to be too permissive. We do not want to end up saying that demonic possession reduces to certain kinds of mental illness; we want to say that theories of mental illness replaced those of demonic possession.
- G. We should also note a connection to one of our themes: In general, the direction of reduction is away from that which is epistemically accessible. Relatively observationally accessible notions, such as temperature, are reduced to relatively inaccessible notions, such as molecular motion. Hence, we see recurring tension between the epistemic modesty emphasized by empiricists and the explanatory, reductive, and metaphysical ambitions that seem to crop up as we look at science's aspirations.

Essential Reading:

Nagel, "Issues in the Logic of Reductive Explanations," in Curd and Cover, *Philosophy of Science: The Central Issues*, pp. 905–921.

Supplementary Reading:

Feyerabend, "How to Be a Good Empiricist—A Plea for Tolerance in Matters Epistemological," in Curd and Cover, *Philosophy of Science: The Central Issues*, pp. 922–949.

Questions to Consider:

1. To a first approximation, proponents of reductionism in a given domain think that a whole is best explained in terms of its parts. Opponents of reductionism in that domain think that the parts can best be understood in terms of their place in the whole. Do you, as a general matter, favor one of these styles of explanation over the other? If so, why?
2. How close an approximation to reduction do you think is needed to vindicate the positivists' notion of scientific progress? At what point do you think we start to see mere replacement rather than reduction?

Lecture Twenty-Four

Reduction and Physicalism

Scope: Reduction at its most dramatic can be seen in claims about whole disciplines. Many philosophers have been tempted by the view that the social sciences reduce to psychology, which reduces to biology, which reduces to chemistry, which reduces to physics. We examine the prospects for this bold proposal and note the impediments placed in its path by multiply realizable entities and properties. Arguably, beliefs and genes are what they are courtesy of their functional role rather than the material of which they are made. This complicates but does not doom the reductionist project. Is science importantly less unified if, say, psychology does not reduce to biology? If the grand reductionist project fails, need that cast doubt on the primacy of the physical?

Outline

- I. Granting, at least for the sake of argument, that one scientific theory sometimes reduces to another, what are the prospects for a version of the unity-of-science program that requires that virtually every scientific theory reduce (ultimately) to basic physics?
 - A. The prototypical vision behind this approach has such sciences as sociology and economics reduce to psychology, which reduces to biology, which reduces to chemistry, which reduces to physics.
 - B. A major idea behind this unity-of-science picture is that nature is, at bottom, homogeneous; it's made out of the same basic stuff, whatever that turns out to be.
 - C. On the other hand, many phenomena in the world seem to involve *emergence*. The properties of table salt are not very like the properties of sodium or of chlorine. Scientific reductions are supposed to explain away the apparent magic of emergence, though we have seen that reductions often involve significant fudging.
- II. *Functional properties* present a major challenge to the reductionist enterprise.
 - A. Thermometers are functionally defined objects. They are picked out by what they do, not what they're made of.
 - B. Such objects need not have any scientifically interesting material properties in common. They are *multiply realizable* in material terms.
 - C. But this makes reduction problematic.
 1. There do not seem to be any *necessary conditions* specifiable in material terms for being a thermometer.
 2. Without necessary conditions, it is hard to see how we are in a position to offer bridge principles linking thermometer talk to material talk. It's hard to see how material discourse will allow us to say what it takes to be a thermometer.
- III. We turn now to a more serious case. Computational psychology concerns the brain's ability to perform deductive inferences. What are the prospects for reducing computational psychology to the physical?
 - A. It seems clear that computational properties are multiply realizable. They can be realized in computers, in our brains, and presumably, in brains different from ours.
 - B. For this reason, it looks unlikely that computational properties can reduce to or be identified with such physical properties as having a particular arrangement of neurons.
 - C. One response available to the reductionist involves taking certain liberties with the idea of a property.
 1. On this view, computational properties reduce to very complicated physical properties (for example, the property of being either a certain kind of microprocessor or a certain kind of brain).
 2. Having one property on the list would be necessary for having a computational property (if you can list all the possible ways of embodying the property), and whichever property one has is sufficient; thus, the bridge principle provides necessary and sufficient conditions for having the computational property.

3. There is nothing intrinsically untoward about disjunctive properties. The property of being a member of Congress can be thought of as the property of being a member of the House of Representatives or of being a member of the Senate.
 4. But do objects that share a property need to be similar in any interesting respect when described at the level of the reducing theory?
 5. Perhaps it suffices for our purposes to note that the giant disjunctive property is not one that a scientist would recognize or be interested in. This is, at best, a philosopher's reduction.
- D. Instead of widening the reducing properties, one could try narrowing the reduced property. It would then be computing-in-humans that gets reduced to the physical.
1. Such an approach does not reduce computing to the physical but only this narrower property. But perhaps this narrower property remains scientifically interesting.
 2. The same worry reasserts itself, however. It is very far from clear that computational properties are physically realized in the same way in all human brains. Brains seem to be able to implement programs in multiple ways.
 3. As this narrowing process continues, we move toward losing everything interesting about the reduction.
- IV. Similar considerations apply to many other properties that figure in science. The property of being money, for instance, will be very difficult to characterize in physical terms. The point is not just that there is a many-to-one relation between physical properties and those of some higher-level science; the point is that the physical properties that fit into such bridge laws do not look like they can figure in laws or explanations.
- V. We've so far focused on the gains and progress that result from scientific reductions (when they happen). But we should not assume that reduction involves gains without losses. Explanatory power can be lost even in favorable cases of reduction.
- A. Even if we confine our attention to cases in which there can be genuine explanations, both at the reducing and the reduced level (thus waiving some of the objections discussed above), we risk the loss of explanatory power. The following example is attributable to Alan Garfinkel.
1. Suppose that an ecologist is tracking a rabbit population that varies more or less inversely with the fox population. We have what looks like a legitimate explanation at the level of ecology: "The rabbit's death was due to a high fox population."
 2. There is a legitimate lower-level explanation along the lines of: "The rabbit's death was due to entering a certain (fox-containing) space at a certain time."
 3. Someone sympathetic to reductionism will claim that explanations of the higher-level type can be replaced by explanations of the lower-level type. Ecological explanations and ecological facts ultimately reduce to lower-level biology and physical facts.
 4. But these explanations seem to account for different facts, namely, why the rabbit was eaten at all versus why the rabbit was eaten when and where it was.
 5. The lower-level explanation thus provides more detail; it explains why the rabbit was eaten when and where it was, rather than just why it was eaten. This might seem to favor reductionism.
 6. But it is plausible to claim that sometimes it is precisely the less specific fact that we want explained. The ecological explanation can account for the (supposed) fact that the rabbit probably would have been eaten even if it had taken a different path.
- B. Even if one were to grant that the project of reducing ecological stuff to biological stuff and, ultimately, to physical stuff looks reasonably promising, it's another thing entirely to claim that ecological explanations can be reduced to any other kind of explanation.
- VI. There seems to be conceptual space for versions of physicalism that are only modestly reductive or that are non-reductive.
- A. One can adopt token physicalism without adopting type physicalism. The idea is that every token of a thermometer, of money, and so on is a physical object. But these types are not physical types. This allows us to say that, in one sense, everything is physical, while denying that everything reduces to the physical.
- B. *Supervenience physicalism* provides a way of insisting that the level of fundamental physics is basic without committing oneself to reductionism. Supervenience physicalism says that any two situations identical in all physical respects would have to be identical in all respects.

1. For example, to say that the mental supervenes on the physical is to say that there can be no mental difference without a physical difference.
 2. This makes room for multiple realizability. Supervenience physicalism allows that there can be a difference in physical properties without, say, a difference in mental properties. It denies only that there can be a difference in mental properties without a difference in physical ones. Thus, there can be more than one physical realization of a mental state.
- C. Supervenience physicalism is much weaker than reduction or identity, but in some cases, it might be too strong. If there could be laws governing a type of particle interaction that will never happen, then the laws do not supervene on the actual physical events.
- D. It turns out to be much trickier than one might have thought to articulate the idea that the world or our theories of it are, in some important sense, unified.

Essential Reading:

Fodor, "Special Sciences," in Boyd, Gasper, and Trout, *The Philosophy of Science*, pp. 429–441.

Garfinkel, "Reductionism," in Boyd, Gasper, and Trout, *The Philosophy of Science*, pp. 443–459.

Supplementary Reading:

Kitcher, "1953 and All That: A Tale of Two Sciences," in Curd and Cover, *Philosophy of Science: The Central Issues*, pp. 971–1003.

Questions to Consider:

1. How permissive a notion of a property do you favor? Do you think there is the property of being an odd number or a middle linebacker or a pumpkin? Why or why not?
2. Do you think that ethics supervenes on the physical? In other words, if two situations were physically identical, would they have to be morally identical?

Glossary

analytic/synthetic: Analytic statements have their truth or falsity determined by the meanings of the terms of which they are composed. “No triangle has four sides” is an example of a (supposed) analytic truth. The truth value of synthetic statements, such as “All copper conducts electricity,” depends, not just on what the statement means, but also on what the world is like. Quine denies that any statements are properly regarded as analytic.

a priori/a posteriori: This distinction concerns how the truth or falsity of a statement can come to be known. A statement is knowable *a priori* if the justification of the statement does not depend on experience. You may, in fact, have learned that $2 + 2 = 4$ through experience (counting apples and oranges and such), but if the justification for this claim is not experiential (if, for instance, the claim is analytically true), then it is knowable *a priori*. Statements not knowable *a priori* are knowable only *a posteriori*, that is, in part on the basis of evidence obtained through experience (though not necessarily one’s own experience).

auxiliary hypotheses: We generally have a sense of which hypothesis we mean to be testing. But no hypothesis has any observational implications all by itself; thus, we must include auxiliary hypotheses in order to derive predictions from the hypothesis under test. Even if the predictions prove false, it is possible that the hypothesis under test is true and that the false prediction should be “blamed” on one of the auxiliary hypotheses.

Bayesianism: Although a range of probability-centered approaches to the theory of evidence and confirmation can be considered Bayesian, orthodox Bayesians interpret probability statements as degrees of belief, and they permit a great deal of “subjectivity” in the assignment of prior probabilities. They require that one update one’s degrees of belief in accordance with Bayes’s Theorem.

bridge law: Bridge laws are crucial to scientific reductions, at least as classically understood. If, as is generally the case, the theory to be reduced contains terms that do not appear in the reducing theory, bridge laws are used to connect the vocabulary of the two theories. “Temperature is mean molecular kinetic energy” is a rough statement of a classic bridge law. Without bridge laws, the reduced theory cannot be logically derived from the reducing theory.

causal model (of explanation): The main successor to the covering-law model of explanation, the causal model says that events are explained by revealing their causes.

cognitive meaning: Cognitively meaningful statements are literally true or false. They are contrasted with sentences that don’t aspire to make true assertions (such as questions, commands, and poetry) and, more important for our purposes, with metaphysical statements that, according to the logical positivists, aspire to cognitive meaningfulness but fail to achieve it.

concept empiricism: This position, exemplified by Hume, asserts that any legitimate concept must be traced back to sources in direct experience. Concepts that cannot be so traced (for example, substance) are not genuinely meaningful.

constructive empiricism: Bas van Fraassen’s constructive empiricism combines an empiricist conception of evidence (according to which all evidence is observational evidence, and the distinction between the observable and the unobservable is of great importance) with an anti-empiricist conception of meaning (scientific theories refer to unobservable reality in much the same way that they refer to observable reality). As a result, van Fraassen maintains that good theories are committed to claims about unobservables but that good scientists need not believe what their theories say about unobservables.

constructivism: This term is sometimes rendered *constructionism*. Generally, this is the idea that (some part of) reality is made rather than found. In the context of our course, this idea gets its start with Kuhn’s suggestion that paradigms help determine a scientist’s world or reality. Constructivists tend to be suspicious of the distinction between experience, theories, and beliefs, on the one hand, and reality, on the other.

context of discovery: The empiricisms of the 17th–19th centuries tried to formulate rules that would lead to the discovery of correct hypotheses. The dominant empiricist views of the 20th century relegated discovery to psychology and sociology; they held that scientific rationality applies only in the context of justification. The distinction between discovery and justification has been under pressure since Kuhn’s work became influential.

context of justification: The positivists and other 20th-century empiricists held that, although no method for generating promising hypotheses is available, once a hypothesis has been generated, a logic or method can be found

by which its justification can be assessed. They thus held that, although there is no logic of scientific discovery, there is one of scientific justification.

contingent/necessary: A contingently true statement is actually true (another way of saying this is to call the statement “true in the actual world”) but could be false (that is, the statement is false in some possible world). Necessary truths hold in all possible worlds. In certain contexts, particular kinds of necessity or contingency are at work. For instance, it is physically necessary (that is, necessary given the laws of nature) that copper conduct electricity, but it isn’t logically necessary that it do so (that is, there is no contradiction involved in the idea of nonconductive copper).

corroboration: This is Popper’s term for theories or hypotheses that have survived serious attempts to refute them. Because Popper insists that corroboration has nothing to do with confirmation, he claims that we have no reason to think corroborated theories more likely to be true than untested ones.

counterfactual: Counterfactual conditionals are expressed in the subjunctive, rather than in the indicative mood. “If my coffee cup were made of copper, it would conduct electricity” is a counterfactual conditional. Counterfactuals can be used to test how robust a statement is, that is, how insensitive the truth of the statement is to actual circumstances.

covering-law model: The centerpiece of logical positivism’s philosophy of explanation, the covering-law model treats explanation as the derivation of the explanandum from an argument containing at least one law of nature.

deduction/deductive logic: This is the relatively unproblematic, well-understood part of logic. It is concerned with the preservation of truth. If an argument is deductively valid, then it is impossible for the premises of the argument to be true while the conclusion is false.

demarcation criterion: A demarcation criterion would provide a basis for distinguishing science from pseudoscience.

determinism: Determinism holds if the state of the universe at a given moment suffices to exclude all outcomes except one. Generally, determinism is understood as causal determinism; the state of the universe at a given moment causally determines the outcome at the next moment. Quantum mechanics suggests that the universe is not deterministic.

disposition: Dispositions manifest themselves only under certain conditions. A substance is soluble (in water) if it is disposed to dissolve when placed in water. Because substances are taken to retain their dispositional properties even when they are not in the relevant circumstances, dispositional properties outrun their manifestations in experience and, thus, pose problems for empiricists.

eliminative reduction: Generally, when some “stuff” (water) or a theory (thermodynamics) is reduced to something else (H₂O or statistical mechanics), the reduced entity or theory does not lose any of its claim to real existence. Sometimes, though, the right sort of reduction eliminates the existence of the thing reduced. When we reduce cases of demonic possession to certain kinds of illness, we thereby show that there were never any cases of demonic possession.

empiricism: A wide range of views can lay claim to this label. They all have in common some conception, according to which experience is the source of some cognitive good (for example, evidence, meaningfulness). See **concept empiricism** and **evidence empiricism**.

entailment: Statement A entails statement B if it is impossible for A to be true without B being true.

epistemology: A fancy Greek word meaning the theory of knowledge and justification.

evidence empiricism: This is the thesis that all of our evidence (at least all of our evidence for synthetic propositions) ultimately derives from experience. Rationalists, in contrast, think that some synthetic statements can be justified on the basis of reason alone.

exemplar: An exemplar is a Kuhnian paradigm in the narrow sense of that term. Exemplars are model solutions to scientific puzzles. Exemplars loom very large in scientific education, according to Kuhn.

explanandum: A fancy Latin word meaning “that which is explained.” It often refers to a *sentence* describing the event (or whatever) being explained.

explanans: A fancy Latin word meaning “that which is explaining.”

explanatory inference: See **inference to the best explanation**.

falsificationism: Popper’s demarcation criterion and his conception of scientific testing are generally combined under this term. Science is distinguished from pseudoscience by the readiness with which scientific claims can be falsified. In addition, scientific testing can falsify but can never confirm theories or hypotheses.

folk psychology: It is much disputed whether folk psychology is a theory or not. We explain one another’s behavior in terms of beliefs, desires, and so on, and this explanatory and predictive practice is folk psychology, whether it merits being considered a psychological theory or not.

functional properties: Some objects are individuated by what they do (or what they’re for), rather than by what they’re made of. Knives can be made of any number of materials; they are united by their purpose or function.

holism: In the context of this course, holism is associated with the work of Quine, who emphasizes holism about testing—no hypothesis can be tested without extensive reliance on auxiliary hypotheses—and holism about meaning—because, in the positivist tradition, testability and meaning are closely linked, statements and terms are meaningful only in the context of a whole theory.

Hume’s fork: Hume’s fork is basically a challenge grounded in his empiricism. All meaningful statements concern either “matters of fact” and are subject to the empirical sciences or “relations of ideas” and are, at bottom, analytic and the proper domain of such disciplines as mathematics. It is a matter of some controversy whether Hume’s fork leaves any space for philosophy.

hypothetico-deductive: Another bit of pure poetry, brought to you courtesy of philosophers of science. The hypothetico-deductive model of confirmation is simple and powerful. It says that a hypothesis is confirmed when true observational consequences can be deduced from it. If the hypothesis (along with auxiliary hypotheses, of course) makes observational predictions that turn out to be false, then the hypothesis is disconfirmed or, perhaps, even refuted.

incommensurability: Literally, this term refers to the lack of a common measure. In the work of Kuhn, Feyerabend, and their successors, incommensurability indicates a range of ways in which competing paradigms resist straightforward comparison. Insofar as two paradigms offer different standards for scientific work and assign different meanings to crucial terms, it will be difficult to assess them in terms of plausibility, promise, and so on.

induction: There is little agreement about how this term should be used. In the narrow sense, induction comprises “more-of-the-same” inferences. A pattern is carried forward to new cases. Some thinkers would assimilate analogical inference to this pattern. In the broad sense, induction includes explanatory inferences, as well as analogical and “more-of-the-same” inferences.

inference to the best explanation: This encompasses a range of inferential practices (such terms as *abductive inference* and *explanatory inference* are sometimes used to mark differences within this range). The general idea is that a theory’s explanatory success provides evidence that the theory is true. This style of argument is crucial to scientific realism but is regarded with some suspicion by empiricists.

instrumentalism: Sometimes, any version of anti-realism about science is called instrumentalist, but it is probably more useful to reserve the term for the idea that scientific theories are tools for predicting observations and, thus, do not have to be true to be good (though they have to lead to true predictions in order to be good).

laws of nature: Not all laws of nature are called laws. Some fundamental and explanatory statements within sciences are called equations, for instance. The philosophical disagreement (between regularity theorists and necessity theorists, mainly) concerns what makes a true, fundamental, and explanatory statement a law of nature.

logical empiricism: See **logical positivism**.

logical positivism: In this course, *logical positivism* and *logical empiricism* are used interchangeably. These terms refer to an ambitious, language-centered version of empiricism that arose in Vienna and Berlin and became the standard view in philosophy of science through the middle of the 20th century. Under the pressure of criticism (largely from within), the positivist program became somewhat more moderate over the years.

metaphysics: This term was generally used pejoratively by the positivists to refer to unscientific inquiries into the nature of reality. These days, most philosophers see room for a philosophical discipline worth calling metaphysics, which addresses such issues as personal identity, the reality of universals, and the nature of causation.

model: Models can be abstract or concrete. In either case, the structure of the model is used to represent the structure of a scientific theory. This semantic approach to theories contrasts with the syntactic approach characteristic of positivism and the “received view” of scientific theories.

naturalism: Naturalism has been enormously influential in recent philosophy. It comes in many flavors, but the central ideas include a modesty about the enterprise of philosophical justification and a consequent emphasis on the continuity between philosophy and science. Naturalists give up on the project of justifying science from the ground up and thereby free themselves to use scientific results (for example, about how perception works) for philosophical purposes.

natural kinds: The contrast is, unsurprisingly, with artificial kinds. The notion of a natural kind can receive stronger and weaker construals. Strongly understood, natural kinds are nature’s joints, grouping things that are objectively similar to one another. Chemical elements might be thought of this way; biological species are a harder case. More weakly, natural kinds are the categories that matter to scientific theorizing.

necessary: See **contingent/necessary**.

necessary condition: A is a necessary condition for B just in case nothing can be B without being A. Being a mammal is a necessary condition for being a whale.

necessitarian view of laws: Unlike regularity theorists, necessitarians maintain that laws of nature do more than just report what invariably happens. Necessitarians think that the laws of nature report relations among universals or similar “deep” features that *make*, for example, copper conduct electricity.

normative: This term contrasts with *descriptive*. Normative claims concern how things ought to be rather than how they are.

objective: A term that probably does more harm than good, but one that is nevertheless nearly impossible to avoid. *Objective* can modify such things as beliefs, in which case it refers to the absence of bias or idiosyncrasy. It can also modify such a term as *existence*, in which case it indicates that something exists independently of its being thought of, believed in, and so on.

ontology: In philosophy, ontology is the part of metaphysics concerned with existence. The ontology of a scientific theory is the “stuff” (objects, properties, and so on) that, according to the theory, exists.

operationalism: Also sometimes called *operationism*, this influential approach to the meaning of scientific terms originated with the physicist P. W. Bridgman. It requires that scientific terms be defined in terms of operations of measurement and detection. This approach is generally thought to be too restrictive.

paradigm: For the narrow use of paradigm, see **exemplar**. In the broad sense, a paradigm includes exemplars but also theories, standards, metaphysical pictures, methods, and whatever else is constitutive of a particular approach to doing science.

partial interpretation: This term contrasts, unsurprisingly, with *full interpretation*. Because the positivists held that meaning arises from experience, they had a difficult time assigning full meaning to statements that go beyond experience. They minimized this problem with their idea of theories as partially interpreted systems. Even if a term such as *fragility* can be applied only to objects that meet certain test conditions, the term is still useful for generating predictions and for connecting observations to one another.

pessimistic induction: This refers to one of the major arguments against scientific realism. Most successful scientific theories have turned out to be false, so we should expect that currently successful theories will turn out to be false.

positivism: In this course, *positivism* is generally used as an abbreviation for *logical positivism*. The term also refers to a 19th-century version of empiricism associated with August Comte. Comte defended a more extreme version of empiricism than did “our” positivists. For instance, he denied that science aspires to explain phenomena.

posterior probability: This is the probability of a hypothesis given some evidence. It is represented as $P(H/E)$ or as $P(H/E \& B)$ if we want to make the role of background evidence explicit. $P(H/E)$ is usually spoken as “the probability of H on E” or “the probability of H given E.”

prior probability: This can either mean the probability of a hypothesis before any evidence at all has been gathered or the probability of the hypothesis before a particular piece of evidence is in. Either way, the prior probability is written $P(H)$.

probability: A mathematical notion, but one that can receive a range of interpretations. We are mainly concerned with Bayesians, for whom probabilities are understood as degrees of belief. Others understand probability statements in terms of (actual or idealized) frequencies or physical propensities, among other possibilities.

problem of old evidence: It would seem that any evidence we already know to be true should receive a probability of 1. But if we plug that value into Bayes’s Theorem, we can see that any evidence that has a probability of 1 cannot confirm any hypothesis in the slightest.

rational reconstruction: Popper and the positivists tended to offer rational reconstructions of scientific practice. A rational reconstruction characterizes the justified core of a practice, rather than the practice as a whole. Largely as a result of Kuhn’s work, philosophers have been less confident in recent years that they can isolate the rational core of science.

realism: See **scientific realism**.

“received view” of theories: See **syntactic conception of theories**.

reduction: A reduction occurs when a more general theory can account for the (approximate) truth of a more specific theory. The standard or classical account of reduction favored by the positivists requires the reduced theory to be derivable from the reducing theory plus suitable bridge laws. Reduction, insofar as it happens, appears to offer an unproblematic sense in which science makes progress.

regularity view of laws: Regularity theorists maintain that laws of nature comprise a subset of nature’s regularities, namely, things that always happen. Laws do not involve any kind of causal necessity, as they do on the rival necessitarian conception of laws.

relativism: The kind of relativism at issue in this course concerns justification or truth. A relativist denies that standards of justification or truth can be applied independently of such things as theories, paradigms, or class interests. The standards are then said to be relative to the theories or interests. Objectivists about justification or truth think that we can make useful sense of these notions independently of our theories or interests.

research program: Lakatos’s notion of a research program is loosely analogous to Kuhn’s notion of a paradigm (in the broad sense). Lakatos allows competition among research programs and imposes a more definite structure on his research programs than Kuhn had on his paradigms. Research programs involve a *hard core* of claims that are not subject to test and a *protective belt* of claims that can be modified in the light of experience.

scientific realism: Another idea that comes in several flavors, scientific realism has at its core the claims that scientific theories aim to correctly depict both unobservable and observable reality and that, in general at least, adopting a scientific theory involves believing what it says about all of reality.

scientific revolution: This term was made famous by Kuhn, but one needn’t be a Kuhnian to think that Newton, Darwin, and Einstein, among others, revolutionized science. Kuhn is skeptical about whether traditional notions of progress and accumulation hold across revolutions, but any view of the history of science will have to make some sense of the enormous changes to scientific practice that have occasionally taken place.

semantic conception of theories: Against the received, or syntactic, view of theories, the semantic approach treats theories as sets of models rather than as axiomatic systems. The semantic approach does not rely as heavily as does the syntactic on the distinction between observable and unobservable reality.

strong program: The strong program is an influential approach within the sociology of science. It seeks to explain scientific behavior by examining the psychological and sociological causes of beliefs and decisions. The strong program’s most controversial component is the symmetry principle, according to which the truth or justification of a belief should play no role in explaining its acceptance.

sufficient condition: A is a sufficient condition for B just in case anything that is A must be B. Being made of copper is sufficient for being metallic.

supervenience: To say that one domain supervenes on another is to say that there can be no change at the “upper” level without a change at the “lower” level. For instance, to say that the domain of the psychological supervenes on the domain of the physical is to say that any two situations that are physically identical would have to be psychologically identical.

syntactic conception of theories: Also known as the received view of theories, this approach conceives theories as systems of sentences modeled, more or less, on geometry. The fundamental laws of the theory are the unproved axioms. In its classic, positivist incarnations, meaning “flows up” into the theory from observation statements, and a theory is, thus, a “partially interpreted formal system.”

synthetic: See **analytic/synthetic**.

teleological explanation: An explanation that makes reference to a purpose is said to be teleological. Such explanations are prevalent in biology (creatures have hearts for the purpose of pumping blood) and psychology (we explain behavior as goal-directed). Philosophers have worked hard to reconcile teleological explanation with non-purposive explanation.

theory-ladenness of observation: Another position that comes in various strengths, claims about theory-ladenness range from uncontroversially modest ones (for example, that the theory one holds will affect how observations are described) to highly controversial ones (notably that observation cannot provide any sort of neutral evidence for deciding between theories or paradigms).

underdetermination: This is generally understood to be shorthand for “underdetermination of theory by evidence.” This thesis is particularly associated with Quine’s holism. For any given set of observations, more than one theory can be shown to be logically compatible with the evidence. More threatening versions of underdetermination maintain that even if additional criteria are imposed (mere logical consistency with the data is, after all, rather weak), no rational basis for settling on a theory will emerge.

unificationist models (of explanation): A recently influential approach, according to which science explains by minimizing the number of principles and argument styles we have to treat as basic. Understanding is increased when the number of unexplained explainers is minimized.

universal generalization: This is the logical form of most laws of nature. “All As are Bs” is the easiest rendering in English of this form.

verification principle: Though it has never quite received a satisfactory formulation, the verification (or verifiability) principle of meaning stood at the center of the logical positivist program. It asserts, roughly, that the meaning of any empirical statement is the method of observationally testing that statement.

Philosophy of Science

Part III

Professor Jeffrey L. Kasser



THE TEACHING COMPANY ®

Jeffrey L. Kasser, Ph.D.

Teaching Assistant Professor, North Carolina State University

Jeff Kasser grew up in southern Georgia and in northwestern Florida. He received his B.A. from Rice University and his M.A. and Ph.D. from the University of Michigan (Ann Arbor). He enjoyed an unusually wide range of teaching opportunities as a graduate student, including teaching philosophy of science to Ph.D. students in Michigan's School of Nursing. Kasser was the first recipient of the John Dewey Award for Excellence in Undergraduate Education, given by the Department of Philosophy at Michigan. While completing his dissertation, he taught (briefly) at Wesleyan University. His first "real" job was at Colby College, where he taught 10 different courses, helped direct the Integrated Studies Program, and received the Charles Bassett Teaching Award in 2003.

Kasser's dissertation concerned Charles S. Peirce's conception of inquiry, and the classical pragmatism of Peirce and William James serves as the focus of much of his research. His essay "Peirce's Supposed Psychologism" won the 1998 essay prize of the Charles S. Peirce Society. He has also published essays on such topics as the ethics of belief and the nature and importance of truth. He is working (all too slowly!) on a number of projects at the intersection of epistemology, philosophy of science, and American pragmatism.

Kasser is married to another philosopher, Katie McShane, so he spends a good bit of time engaged in extracurricular argumentation. When he is not committing philosophy (and sometimes when he is), Kasser enjoys indulging his passion for jazz and blues. He would like to thank the many teachers and colleagues from whom he has learned about teaching philosophy, and he is especially grateful for the instruction in philosophy of science he has received from Baruch Brody, Richard Grandy, James Joyce, Larry Sklar, and Peter Railton. He has also benefited from discussing philosophy of science with Richard Schoonhoven, Daniel Cohen, John Carroll, and Doug Jesseph. His deepest gratitude, of course, goes to Katie McShane.

Table of Contents

Philosophy of Science

Part III

| | |
|---|---------|
| Professor Biography | i |
| Course Scope | 1 |
| Lecture Twenty-Five New Views of Meaning and Reference | 3 |
| Lecture Twenty-Six Scientific Realism | 6 |
| Lecture Twenty-Seven Success, Experience, and Explanation..... | 9 |
| Lecture Twenty-Eight Realism and Naturalism..... | 12 |
| Lecture Twenty-Nine Values and Objectivity | 14 |
| Lecture Thirty Probability | 17 |
| Lecture Thirty-One Bayesianism..... | 20 |
| Lecture Thirty-Two Problems with Bayesianism..... | 23 |
| Lecture Thirty-Three Entropy and Explanation | 26 |
| Lecture Thirty-Four Species and Reality | 29 |
| Lecture Thirty-Five The Elimination of Persons?..... | 32 |
| Lecture Thirty-Six Philosophy and Science | 35 |
| Timeline | Part I |
| Glossary | Part II |
| Biographical Notes | 37 |
| Bibliography | 40 |

Philosophy of Science

Scope:

With luck, we'll have informed and articulate opinions about philosophy and about science by the end of this course. We can't be terribly clear and rigorous prior to beginning our investigation, so it's good that we don't need to be. All we need is some confidence that there is something about science special enough to make it worth philosophizing about and some confidence that philosophy will have something valuable to tell us about science. The first assumption needs little defense; most of us, most of the time, place a distinctive trust in science. This is evidenced by our attitudes toward technology and by such notions as who counts as an expert witness or commentator. Yet we're at least dimly aware that history shows that many scientific theories (indeed, almost all of them, at least by one standard of counting) have been shown to be mistaken. Though it takes little argument to show that science repays reflection, it takes more to show that philosophy provides the right tools for reflecting on science. Does science need some kind of philosophical grounding? It seems to be doing fairly well without much help from us. At the other extreme, one might well think that science occupies the entire realm of "fact," leaving philosophy with nothing but "values" to think about (such as ethical issues surrounding cloning). Though the place of philosophy in a broadly scientific worldview will be one theme of the course, I offer a preliminary argument in the first lecture for a position between these extremes.

Although plenty of good philosophy of science was done prior to the 20th century, nearly all of today's philosophy of science is carried out in terms of a vocabulary and problematic inherited from logical positivism (also known as logical empiricism). Thus, our course will be, in certain straightforward respects, historical; it's about the rise and (partial, at least) fall of logical empiricism. But we can't proceed purely historically, largely because logical positivism, like most interesting philosophical views, can't easily be understood without frequent pauses for critical assessment. Accordingly, we will work through two stories about the origins, doctrines, and criticisms of the logical empiricist project. The first centers on notions of meaning and evidence and leads from the positivists through the work of Thomas Kuhn to various kinds of social constructivism and postmodernism. The second story begins from the notion of explanation and culminates in versions of naturalism and scientific realism. I freely grant that the separation of these stories is somewhat artificial, but each tale stands tolerably well on its own, and it will prove helpful to look at similar issues from distinct but complementary angles. These narratives are sketched in more detail in what follows.

We begin, not with logical positivism, but with a closely related issue originating in the same place and time, namely, early-20th-century Vienna. Karl Popper's provocative solution to the problem of distinguishing science from pseudoscience, according to which good scientific theories are *not* those that are highly confirmed by observational evidence, provides this starting point. Popper was trying to capture the difference he thought he saw between the work of Albert Einstein, on the one hand, and that of such thinkers as Sigmund Freud, on the other. In this way, his problem also serves to introduce us to the heady cultural mix from which our story begins.

Working our way to the positivists' solution to this problem of demarcation will require us to confront profound issues, raised and explored by John Locke, George Berkeley, and David Hume but made newly urgent by Einstein, about how sensory experience might constitute, enrich, and constrain our conceptual resources. For the positivists, science exhausts the realm of fact-stating discourse; attempts to state extra-scientific facts amount to metaphysical discourse, which is not so much false as meaningless. We watch them struggle to reconcile their empiricism, the doctrine (roughly) that all our evidence for factual claims comes from sense experience, with the idea that scientific theories, with all their references to quarks and similarly unobservable entities, are meaningful and (sometimes) well supported.

Kuhn's historically driven approach to philosophy of science offers an importantly different picture of the enterprise. The logical empiricists took themselves to be explicating the "rational core" of science, which they assumed fit reasonably well with actual scientific practice. Kuhn held that actual scientific work is, in some important sense, much less rational than the positivists realized; it is driven less by data and more by scientists' attachment to their theories than was traditionally thought. Kuhn suggests that science can only be understood "warts and all," and he thereby faces his own fundamental tension: Can an understanding of what is intellectually special about science be reconciled with an understanding of actual scientific practice? Kuhn's successors in sociology and philosophy wrestle (very differently) with this problem.

The laudable empiricism of the positivists also makes it difficult for them to make sense of causation, scientific explanation, laws of nature, and scientific progress. Each of these notions depends on a kind of connection or

structure that is not present in experience. The positivists' struggle with these notions provides the occasion for our second narrative, which proceeds through new developments in meaning and toward scientific realism, a view that seems as commonsensical as empiricism but stands in a deep (though perhaps not irresolvable) tension with the latter position. Realism (roughly) asserts that scientific theories can and sometimes do provide an accurate picture of reality, including unobservable reality. Whereas constructivists appeal to the theory-dependence of observation to show that we help constitute reality, realists argue from similar premises to the conclusion that we can track an independent reality. Many realists unabashedly use science to defend science, and we examine the legitimacy of this naturalistic argumentative strategy. A scientific examination of science raises questions about the role of values in the scientific enterprise and how they might contribute to, as well as detract from, scientific decision-making. We close with a survey of contemporary application of probability and statistics to philosophical problems, followed by a sketch of some recent developments in the philosophy of physics, biology, and psychology.

In the last lecture, we finish bringing our two narratives together, and we bring some of our themes to bear on one another. We wrestle with the ways in which science simultaneously demands caution and requires boldness. We explore the tensions among the intellectual virtues internal to science, wonder at its apparent ability to balance these competing virtues, and ask how, if at all, it could do an even better job. And we think about how these lessons can be deployed in extra-scientific contexts. At the end of the day, this will turn out to have been a course in conceptual resource management.

Lecture Twenty-Five

New Views of Meaning and Reference

Scope: A new philosophical theory of reference and meaning makes it easier to face problems of incommensurability; philosophers can now more readily say that we have a new theory about the same old mass rather than a theory of Einsteinian mass competing with a theory of Newtonian mass. The new theory, for better and for worse, also makes it easier to talk about unobservable reality. In this lecture, we explore this new approach to meaning and reference, along with a new conception of scientific theories that accompany it. Scientific theories are now sometimes conceived in terms of models and analogies, rather than as deductive systems. We also consider some legitimate worries the once-received view poses for the new view.

Outline

- I. At this point, we begin bringing our two narratives together by integrating issues of meaning and reference into our recent discussions of explanation and allied notions. We have been tacitly relying on a fairly standard philosophical account of reference, according to which we typically pick things out by correctly describing them.
 - A. Meaning and reference are distinct. *Albert Einstein* and *the discoverer of special relativity* co-refer, but they do not have the same meaning. Likewise, *creature with a heart* applies to all the same things as *creature with a kidney*, but they don't mean the same thing.
 - B. In a standard understanding, a description such as *the favorite physicist of the logical positivists* must correctly pick out a unique individual (for example, Einstein) in order to refer.
 - C. Suppose that, unbeknownst to me, Werner Heisenberg turns out to be the favorite physicist of the logical positivists. In that case, I may think I am using the phrase to refer to Einstein, but I am really referring to Heisenberg.
 - D. As we have seen, the logical positivists treated meaning and reference as relatively unproblematic for observational terms and as quite problematic for theoretical terms.
 - E. A common version of this approach does not provide reference for theoretical terms at all; the parts of scientific theory that are not about experience do not directly refer to the world and do not aspire to truth. Talk of quarks just serves to systematize and predict observation.
 - F. Less stringent empiricists allowed theoretical terms to refer and treated them in the standard way. This is the approach taken by Thomas Kuhn.
 1. For Kuhn, reference is fairly easy to secure, because a term refers only to the world-as-described-by-the-paradigm. Thus, in Kuhn's view, such a term as *phlogiston* refers just as surely as *oxygen* does to something that can cause combustion; both refer to crucial causes of combustion, as identified by their paradigms.
 2. This makes reference too easy to secure. Most philosophers find it much more natural to say that phlogiston never existed, and the term *phlogiston* never referred to anything.
 - G. On the other hand, the standard view makes reference too hard to secure. If Benjamin Franklin misdescribes electricity, then, because there is nothing meeting his description, he is not talking about electricity at all.
 - H. Similarly, this descriptive conception of reference looms large in the somewhat exaggerated incommensurability arguments of Kuhn and Paul Feyerabend. If enough descriptive content changes, the reference will likely change with it. Thus, when descriptions of mass change across theories, the new theory often refers to something new, namely, mass-as-conceived-by-the-theory. For this reason, Einstein cannot offer a better theory of the same mass as Newton's, and this makes progress and accumulation difficult.
- II. A new conception of reference emerged (mainly in the 1970s) that makes it easier to talk about unobservable reality and to keep talking about the same things or properties, even across major scientific changes. On this

view, reference (for certain kinds of terms) is secured through a historical chain, rather than through a description. It is often called a *causal theory of reference*.

- A. Proper names provide the easiest starting point. If you say, “James Buchanan, the 14th president” (he was actually the 15th), you are still referring to Buchanan.
 - 1. Buchanan’s name was attached to him via a kind of baptismal event, not a description. This is a stipulation.
 - 2. My use of his name is linked to previous uses in a causal chain that terminates in the baptismal event. I intend to refer to the same man as the person from whom I learned the name, and so on, back through the chain to the first link.
 - B. Similar things can be said of “natural-kind” terms, such as biological species. We would like a theory that allows us to say that people who thought that whales were fish nevertheless referred to whales.
 - 1. The reference of such terms gets fixed via an archetypal specimen: Whales are creatures like this one.
 - 2. *Like this one* means having the same deep or essential properties. For chemical elements, it will be their atomic numbers.
 - C. There is a division of linguistic labor involved in this picture. I do not have to know much about James Buchanan in order to talk about him. Similarly, I do not have to know deep facts about whales in order to succeed in talking about them.
 - D. This new conception of reference had an unexpected consequence: It helped make metaphysical discourse look more respectable than it had to the positivists.
 - 1. If Hesperus and Phosphorous are two different names (rather than descriptions) for the planet Venus, then it is necessarily true that Hesperus is Phosphorous, and this is not a necessity that is analytic and knowable *a priori*. Room is made for a notion of metaphysical necessity that does not reduce to conceptual necessity.
 - 2. This talk of a deep structure shared by all members of natural kinds, such as chemical elements, also rehabilitates, to a significant extent, the notion of essences, which had long been thought unduly metaphysical. These deep structural properties look scientifically respectable.
 - E. This approach to reference also makes incommensurability look much less threatening than it had. Insofar as this approach can be made to work, theory change, even across revolutions, can involve competing theories about the same “stuff,” rather than just theories about different “stuff.”
 - F. The causal/historical approach does make it easier to talk about unobservable reality in a meaningful way. On the assumption that water has a deep structure responsible for its nature, the historical chain approach allows one to talk meaningfully about that structure.
 - G. However, we can never encounter specimens of the purported objects of some theoretical terms. We cannot point at an electron and say, “I mean to be talking about everything that is like that thing.” Given how messy the notion of causation is and how messy the causal chain would have to be, it would be hard to pick out an electron as what is responsible for the streak in the cloud chamber.
 - H. The historical chain approach can also make it too easy to refer to unobservable reality. We don’t want to count someone as referring to oxygen when using the term *phlogiston*, even though oxygen is what is causally responsible for combustion.
- III. A new conception of scientific theories also makes it easier to extend meaning and reference to unobservable reality.
- A. The received view of theories treats them as deductive systems, which get interpreted when some terms are explained experientially. Statements involving theoretical terms generally receive only a partial interpretation.
 - B. A newer conception of theories draws on the notion of a model.
 - 1. A model can be formal. For instance, a wave equation can be used to model waves of sound, or of light, and so on.
 - 2. Models can also be material, in which case they interpret the theory in terms of real or imaginary objects, rather than abstract structures. For example, gas molecules are modeled as small, solid balls.
 - C. Logical positivism assigns only a modest role to models.

1. Models can serve a heuristic function. They involve pictures or analogies that are useful for understanding a theory or for using it.
 2. But the model is not part of the theory, and the theory, not the model, is what says what the phenomena in its domain are like.
- D. But if the model continues to be useful in enough different contexts, it becomes more than just an aid or a supplement to the real theory. A good enough model virtually becomes the theory. Models loom large in scientific practice.
- E. The *semantic conception of theories* identifies a theory with the entire class of its models. A correct theory will have the real world as one of its models. An ecological theory can be interpreted, for example, via patterns of shapes and colors on a computer screen, or via mathematical equations, or via actual patterns of fox and rabbit populations.
1. The big departure from the received view is that semantic approaches allow theoretical terms to be interpreted directly through models, rather than requiring that interpretation always arise through observation.
 2. The semantic conception thus allows a role for analogical and metaphorical reasoning in science. These types of reasoning can provide literal content to what our theory says about unobservable reality.
- F. But how do we restrict the permitted types of modeling and analogical reasoning?
1. What stops someone from claiming to understand absolute simultaneity on the model of local simultaneity?
 2. With some theories, most notably quantum mechanics, there seems to be powerful reasons to resist taking models too seriously.

Essential Reading:

Putnam, "Explanation and Reference," in Boyd, Gasper, and Trout, *The Philosophy of Science*, pp. 171–185.

Kitcher, "Theories, Theorists and Conceptual Change," in Balashov and Rosenberg, *Philosophy of Science: Contemporary Readings*, pp. 163–189.

Supplementary Reading:

Spector, "Models and Theories," in Brody and Grandy, *Readings in the Philosophy of Science*, pp. 44–57.

Questions to Consider:

1. Do you think that science should strive to be as free of metaphor and analogy as possible? Why or why not?
2. Suppose there were a substance that behaved just like water (for example, we could drink it) but had a quite different molecular structure. Would that substance count as water? Why or why not?

Lecture Twenty-Six

Scientific Realism

Scope: The semantic developments sketched in the previous lecture make room for the doctrine of *scientific realism*, which requires that science “talk about” unobservable reality in much the same way that it talks about observable reality. In this lecture, we examine the varieties and ambitions of scientific realism, contrast it with empiricism and constructivism, and confront two major challenges to realist interpretations of science.

Outline

- I. A number of considerations convinced many philosophers that there is no interesting distinction to be drawn between observational and theoretical language. Without such a distinction, logical positivism is more or less dead. The epistemology of empiricism can live on, but it will have to take a different form (as we will see).
 - A. The new conceptions of meaning and reference that we canvassed in the last lecture suggested that our semantic reach can extend farther beyond observation than the positivists had thought.
 - B. A relatively modest descendant of a point made by Kuhn and Feyerabend also contributed to the new skepticism about the observational/theoretical distinction. They insisted that theories shape what we see and how we describe what we see.
 1. Most philosophers were not enormously impressed by the argument that our theories “infect” our observations. By and large, philosophers accepted only modest versions of this claim.
 2. But they did become convinced that our theories “infect” our observational *language*. We use theoretical terms (such as *radio*) to talk about observable things. Such talk is fully, not partially, meaningful. The majority of philosophers gave up on the idea that anything worth calling science could be done in a language that was sanitized of reference to unobservable reality.
 - C. Conversely, we can use observation terms to describe unobservable objects (as when we picture gas molecules as little billiard balls).
 - D. Thus, the distinction between observable and theoretical language does not line up with the distinction between observable and unobservable objects.
- II. Statements about unobservable reality, then, can be true or false in the same way that statements about observable reality can. This makes room for *scientific realism*, a view that requires that science aim at accurately depicting unobservable as well as observable reality. What else is involved in scientific realism?
 - A. Metaphysical modesty is a requirement: The way the world is does not depend on what we think about it.
 - B. Epistemic presumptuousness is also a requirement: We can come to know the world more or less as it is.
 - C. Although each of these theses holds considerable appeal, they tend to work against each other. The more independent the world is of us and our thought, the more pessimistic it seems we should be about our prospects for knowing it.
- III. We have seen two anti-realist positions that reject metaphysical modesty, and these can be compared with two realist positions that accept different versions of metaphysical modesty.
 - A. The logical positivists reject questions about the way the world is. They consider such questions invitations to metaphysics.
 - B. For Kuhnian and other constructivists, the way the world is *does* depend on what we think about it.
 - C. For “hard” realists, *the way the world is* means that some distinctions, similarities, and kinds are, as it were, “out there.” The world determines that gold is a real kind, all the instances of which share important properties, while *jade* names an unreal kind, two different kinds of things (jadeite and nephrite) that go by one name.
 - D. For “soft” realists, *the way the world is* means only that, given certain interests and aptitudes, it makes good sense to categorize things in one way rather than another (for example, to think of gold as one kind of thing, but jade as two). Our best theories take our interests into account, but they are still responsible to a mind-independent world.

- E. Hard realists think that the job of science is to find out the way the world truly is, and this goal has nothing to do with contingent human limitations. Soft realists think that the aim of science is to organize a mind-independent world in one of the ways that makes most sense to us. Soft realists generally permit the idea that incompatible theories could be equally good, while this is much harder to grant according to hard realism.
 - F. Hard realism runs the danger of being too restrictive, while soft realism can easily become too permissive. As we will see in later lectures, it's not clear that the world has many kinds that live up to hard realist standards. On the other hand, not every classification scheme that's good for certain purposes thereby gets to claim that the classification is correct.
- IV. Turning from metaphysical issues of modesty to epistemological issues of presumptuousness, we can review some previously examined positions and compare them to a couple of versions of scientific realism.
- A. Logical positivists think that we cannot get evidence that bears on the truth of statements about unobservable reality. Therefore, we should not presume to have knowledge that so thoroughly outruns the evidence.
 - B. For Karl Popper, it is possible that we could come to know the world as it is, but because there is no usable notion of confirmation, we'll never be in a position to claim such knowledge about anything.
 - C. For Kuhn and other constructivists, knowledge of the way the world is would require stepping out of our intellectual and perceptual skins. Even if the project made metaphysical and semantic sense, it would be excessively epistemically presumptuous.
 - D. For "optimistic" realists, our best scientific theories provide knowledge of the way the world is (including unobservable reality). This is the most epistemologically presumptuous view out there, but it's not a crazy or uncommon one. However, this view sets things up so that if major scientific theories are false, then scientific realism is false, and that seems undesirable.
 - E. For "modest" realists, it is reasonable to hope that science can, and sometimes does, provide knowledge of the way the world is. Such thinkers count as realists because they think science has a reasonable chance of getting the world right, but they need not think that it has done so.
- V. The most important debates among realists and between realists and their opponents have concerned epistemic issues: How confident should we be that science does, or at least can, provide us with knowledge of unobservable reality?
- A. The *underdetermination of theory by data* made its first appearance in W. V. Quine's work.
 - 1. It is often the case that all the currently available evidence fails to decide between two competing theories. But this needn't trouble the realist much so long as science has some decent prospect of determining which theory is true.
 - 2. Stronger versions of underdetermination claim that all possible evidence underdetermines theory choice. This is awkward for the realist, who needs to claim that (at most) one of the theories is true.
 - B. A couple of replies are available to the realist.
 - 1. One is to deny that we can always find genuine theories that compete with a given theory. For example, I would not be proposing a new theory if I switch the terms *positive* and *negative* so that electrons have a positive charge and protons have a negative charge. This is the same theory in a verbally incompatible form.
 - 2. Realists can also appeal to principles governing the way to run a web of belief and claim that of two theories that fit the data equally well, one might, nevertheless, receive more evidential support than the other.
 - C. The other major obstacle to realism is an important historical argument called the *pessimistic induction*.
 - 1. We can find cases from the history of science of theories that did as well or better than current theories by the best evidential standards of the day. Because we now know those theories to be false, we should not think our best theories likely to be true.
 - 2. This objection follows Kuhn in thinking that the history of science is our best guide to how science should be done. But it tries to demonstrate that history shows that realism is unwarranted, because the best standards of actual science permit false theories to thrive.
 - D. The realist has room to maneuver here, as well.

1. If some version of the traditional approach to scientific reduction can be defended, then one can claim that superseded theories are preserved by being reduced into superseding theories.
2. Realism might need to narrow its ambitions and claim only that parts of our best theories are likely to be true or that only some of our best theories are likely to be true. Not all aspects of our theories are equally accessible to us or equally well tested.
3. Realism could be defended concerning the mathematical structures involved in our best theories, rather than the entities posited by them. Nicolas Carnot worked out many of the basic ideas of thermodynamics, despite the fact that he mistakenly thought of heat as a kind of fluid.

Essential Reading:

Nagel, "The Cognitive Status of Theories," in Balashov and Rosenberg, *Philosophy of Science: Contemporary Readings*, pp. 197–210.

Laudan, "A Confutation of Convergent Realism," in Balashov and Rosenberg, *Philosophy of Science: Contemporary Readings*, pp. 211–233 (also in Boyd, Gasper, and Trout, *The Philosophy of Science*, pp. 223–245, and in Curd and Cover, *Philosophy of Science: The Central Issues*, pp. 1114–1135).

Supplementary Reading:

Psillos, "The Present State of the Scientific Realism Debate," in Clark and Hawley, *Philosophy of Science Today*.

Questions to Consider:

1. How sympathetic are you to the idea that science does (or at least can) "carve nature at its joints"? What considerations could help you decide between a hard realism like this and a soft realism or an anti-realism?
2. How independent of thought does the notion of "the world" or "the truth" seem to you? Surely my thinking something doesn't make it so. But what about the idea that any statement that would be agreed upon "at the end of inquiry" would have to be true? Is this conception of truth too metaphysically immodest? Why or why not?

Lecture Twenty-Seven

Success, Experience, and Explanation

Scope: Realists defend their position as the best explanation for the success of science. Anti-realists point to a number of successful-but-false theories in the history of science. Under what conditions, if any, does the success of a theory give us good reason to think that it is true (including in what it says about unobservable reality)? We consider empiricist arguments that the demand for an explanation of the success of science begs the question against anti-realism and constitutes an invitation to metaphysics. We also contrast scientific realism with Bas van Fraassen's constructive empiricism, which combines the semantic claims of realism with the suggestion that scientists shouldn't believe what their theories say about unobservable reality.

Outline

- I. *Inference to the best explanation* is the main style of argument for inferring from observable phenomena to unobservable phenomena.
 - A. The straightforward argument for realism is often called the “no miracles” argument. The natural sciences have been tremendously successful, and a fairly strong version of realism (the claim that our best scientific theories are at least approximately true, including what they say about unobservable reality) provides the best explanation for this striking fact.
 1. Two kinds of success matter to the “no miracles” argument: predictive and technological.
 2. Given that some kinds of predictive and technological success are cheap, the “no miracles” argument has got to set the bar pretty high if it is to claim that it would be a miracle that science could do what it has done without its theories being at least approximately true.
 3. Even so, the argument runs up against the pessimistic induction argument discussed in the preceding lecture. Predictively accurate and technologically fruitful theories from the past have been shown to be false. Other generations would have been just as entitled to use the “no miracles” argument, but they would have been wrong; thus, we should not help ourselves to this argument.
 - B. The realist needs to require *novel* predictive success before a theory can justifiably be considered true. If a theory explains only data that are already “in,” a competing explanation is available for the theory's success, i.e., that it was designed to accommodate the data. Novel predictions preclude this explanation and thereby favor the explanation that the theory works because it is true.
 1. Novelty is tricky to characterize. It's neither a straightforwardly temporal nor a straightforwardly psychological notion.
 2. Even if we confine ourselves to novel predictions, the “no miracles” argument is not unproblematic. The wave theory of light generated precise, surprising, and correct predictions, but it is false.
 3. Another response available to the realist is to argue that the success of prediction is due to a part of the wave theory that was, in fact, correct and that error does not disqualify part of the theory from being true. Only for the highly tested parts of the theory will the realist's explanation of success seem like the best one, and even then, one should admit that it is fallible.
- II. Empiricists challenge the whole appeal to inference to the best explanation in the first place. They ask whether the success of scientific theories needs to be explained at all and whether positing the truth of what scientific theories say about unobservables is really the best explanation.
 - A. Van Fraassen uses an evolutionary analogy to resist realism. Theories that generate false predictions tend to get discarded, so it comes as no surprise that the theories that remain generate primarily true predictions. But can this deflationary explanation handle the novel predictive successes of science?
 - B. Many empiricists consider inference to the best explanation questionable when used within science and even more questionable when used about science. Do we have good reason to think that the world will uphold our explanatory ambitions? Do we have good reason to consider explanatory loveliness a mark of truth?
 - C. The status of inference to the best explanation is, thus, quite controversial. Realists argue that such inferences are part and parcel of ordinary and scientific rationality, while empiricists emphasize the

problems with such inferences and claim that unrestricted demands for explanation tend to lead to metaphysical speculation.

III. Van Fraassen's *constructive empiricism* offers a major empiricist alternative to realism.

- A. Van Fraassen agrees with realists about semantic issues: Scientific theories posit observables in fully meaningful ways. Our theories are committed to the existence of such things as electrons.
- B. But it does not follow that we are or should be committed to the existence of electrons.
 - 1. All our evidence is observational evidence, and we shouldn't consider ourselves in a position to attain knowledge of unobservable reality.
 - 2. Thus, we shouldn't believe what our theories say about unobservable reality; at best, we should believe our theories to be *empirically adequate*.
 - 3. While denying that the distinction between observational and theoretical language can do any philosophical work, van Fraassen maintains there is an important difference between observable objects and unobservable ones. From the viewpoint of science, human beings are a certain kind of measuring device, and our evidence is tied to our size, our senses, and so on.
 - 4. Van Fraassen permits inductive arguments from observed phenomena to other observable phenomena, and he permits explanatory inferences to observables. It is inference to unobservables (which induction by itself will not get you) about which he is skeptical.
 - 5. Though van Fraassen does not think that scientists should believe everything their theories say, he does think that they should act as if well-supported theories are true and should use theories for such purposes as experimental design. We can let ourselves be guided by pictures without believing the pictures.
- C. Van Fraassen has shown that the demise of positivism does not mean that empiricism about scientific theories is doomed. But his position is subject to a number of questions.
 - 1. Can the observable/unobservable distinction bear the weight that van Fraassen requires of it? The realist can argue that, despite the fact that we can check only what a theory says about observable reality, we can take methods that we know are reliable with respect to observable reality and apply them to unobservable reality.
 - 2. When we ask why a theory that posits unobservables is predictively accurate and technologically useful, van Fraassen says it is because the theory is empirically adequate. This explanation is likelier than the realist's explanation, but it is very unlovely. Van Fraassen thinks it is no part of science to explain the success of science, but many thinkers find such an explanatory project well motivated.
 - 3. Finally, we can raise some questions about the balance of epistemic modesty and presumptuousness struck by van Fraassen. If we are to be cautious about venturing beyond the observable, why should we not be comparably cautious about venturing beyond the observed? Believing our theories empirically adequate goes enormously beyond the evidence, as the problem of induction shows.

Essential Reading:

Boyd, "On the Current Status of Scientific Realism," in Boyd, Gasper, and Trout, *The Philosophy of Science*, pp.195–222.

Van Fraassen, "Arguments Concerning Scientific Realism," in Curd and Cover, *Philosophy of Science: The Central Issues*, pp. 1064–1087.

Supplementary Reading:

Brown, "Explaining the Success of Science," in Curd and Cover, *Philosophy of Science: The Central Issues*, pp. 1136–1152.

Musgrave, "Realism versus Constructive Empiricism," in Curd and Cover, *Philosophy of Science: The Central Issues*, pp. 1088–1113.

Questions to Consider:

1. Do you think it matters whether one believes a scientific theory or merely accepts it? Does it matter whether one believes some religious doctrine or merely accepts it? Why or why not?
2. To what extent do the major scientific innovations of the last century or so (relativity, quantum mechanics, molecular biology, the rise of psychology, and so on) make scientific realism either harder or easier to defend?

Lecture Twenty-Eight

Realism and Naturalism

Scope: These days, scientific realism is generally offered as the best *scientific* explanation of the success of (some) scientific theories. But many empiricists and constructivists object that this amounts to invoking science to testify on its own behalf. What exactly is the claim of circularity here and how damaging is it? Defenders of *naturalistic epistemology* defend a relatively modest conception of justification and emphasize the continuity of philosophy with the sciences. Radical naturalistic epistemologists, such as Quine, have proposed *replacing* epistemology with scientific psychology. We examine moderate and radical philosophical naturalisms and return to the justification of induction as a test case for naturalized epistemologies. We close by asking whether the naturalistic examination of science looks like it will vindicate or disappoint our hopes about scientific reasonableness.

Outline

- I. The realist asserts and the empiricist denies that inference to the best explanation can make statements about unobservable reality belief-worthy. In the face of this impasse, many realists have adopted an interesting line of partial retreat. They argue that realism is best defended from within a *naturalistic* approach to philosophy.
 - A. Naturalism abandons the project of providing a philosophical justification for science. It gives up on the old, grand conception of philosophy, according to which philosophy can attain *a priori* knowledge through reason alone. But it also gives up on the logical positivists' conception of philosophy as one that tries to achieve valuable results through conceptual analysis alone.
 - B. Naturalism is characterized by the rejection of an extra-scientific standpoint from which science can be assessed. For a naturalist, philosophy and science are continuous with one another.
- II. A naturalistic approach to realism puts scientific realism forward as the best scientific explanation for the success of science. It no longer attempts a philosophical justification of inference to unobservables.
 - A. Scientific realism becomes an empirical hypothesis rather than a philosophical thesis. A naturalized scientific realism takes a scientific look at science and asks whether the successes of science are capable of receiving a scientific explanation. It claims that realism provides the best scientific explanation for the success of science.
 - B. The justification offered for realism is that it meets the standards for explanatory inferences that figure in science itself. No attempt is made to address philosophical worries about whether the standards used in science are legitimate.
 - C. Like the sociology of science, naturalism involves taking a scientific look at science itself. It involves a scientific examination of the conditions under which scientific practices seem reliable. Naturalists who are realists think that this scientific examination turns out differently than the sociologists believe. They think that the methods of current science can be shown to be reasonably reliable.
- III. Two major worries about naturalism arise almost immediately.
 - A. Isn't naturalism troublingly circular? Doesn't it amount to judging science by its own standards?
 - 1. Naturalists are influenced by Kuhn, who suggests that we have no better way of figuring out how science ought to be done than by looking at how it is done.
 - 2. They are also influenced by Quine's holism, according to which no part of the web of belief stands apart from the rest. In such a picture, there will be no distinctively philosophical or distinctively secure knowledge about how to inquire.
 - 3. If one were using science to defend the epistemic credentials of science, then the charge of circularity would be well founded. However, that is not what the naturalists are doing. They repudiate the project of justifying science's epistemic credentials in the first place.
 - 4. Rejecting the demand for a philosophical justification must not be confused with having answered it. Science cannot be vindicated by appealing to science. Naturalism refuses to worry about vindicating science.

- B. Doesn't naturalism threaten to turn philosophy into some mix of biology and psychology, that is, into the scientific study of how perception, inference, and so on happen? And, in doing so, doesn't it lose sight of the distinction between descriptive and normative questions, between "ises" and "oughts"?
 - 1. Though he later moderated his position, Quine initially defended a strong naturalism along just these lines. He suggested that philosophers get out of the knowledge business. Epistemology should become the study of how science generates such ambitious theories on the basis of such slender inputs.
 - 2. Later philosophers in the naturalistic tradition have been less reductive than Quine was. They think that philosophy can use science to help answer philosophical questions without philosophy thereby becoming part of science. The work of philosophers remains primarily conceptual, but it draws on empirical results.
 - C. A naturalistic approach to Nelson Goodman's new riddle of induction can serve as an illustration.
 - 1. The naturalist will try to solve questions about legitimate predicates empirically, not conceptually. We should use our best scientific theories to figure out which predicates are legitimately employed in inductive arguments. If the best explanation for the success of a theory is that it employs the right categories, we have some reason to rely on that theory.
 - 2. This approach does not try to address the big epistemological questions about induction. It assumes that such questions have received favorable answers, and it uses science to help answer smaller problems, such as that of figuring out which inductions are better than others.
 - 3. The anti-naturalist will point to the circularity involved in this defense, while the naturalist will ask how we are supposed to justify anything interesting without using our best theories of the world.
- IV. The naturalistic approach does not automatically vindicate current science. Naturalism can threaten, as well as support, our confidence that current science is reliable.
- A. In some respects, naturalism makes the problem of induction even harder to solve, because it raises the problem of obtaining an adequate description of our inductive practices, and that task is very difficult.
 - B. Many studies in social psychology appear to show that humans reason badly in certain systematic ways. They violate the basic norms of logic and probability theory.
 - C. Evolutionary psychology and evolutionary epistemology suggest that we might be "wired" for some false beliefs about fundamental physics (for example, the impetus theory and a Euclidean geometry of space).
 - D. Many sociologists of science think that their empirical work deflates certain myths concerning the rationality and objectivity widely thought to be characteristic of science. A naturalistic approach to science, they think, is incapable of vindicating something like scientific realism.

Essential Reading:

Quine, "Natural Kinds," in Boyd, Gasper, and Trout, *The Philosophy of Science*, pp. 159–170.
 Godfrey-Smith, *Theory and Reality: An Introduction to the Philosophy of Science*, chapter 10.

Supplementary Reading:

Kornblith, *Naturalizing Epistemology*.

Questions to Consider:

- 1. Naturalized epistemology abandons the project of convincing skeptics that science is justified. Do you think that there are many real-life skeptics about scientific justification? How important do you think it is to respond to such skeptics?
- 2. Can evolutionary epistemology help explain why so much of fundamental physics seems deeply weird to us? Should an evolutionary understanding of human beings alter our conception of what counts as a satisfying explanation, either in physics or in other fields?

Lecture Twenty-Nine

Values and Objectivity

Scope: Recent work in naturalistic epistemology has turned to the social structure of science. This work has been much friendlier to traditional ideals of objectivity than was the strong program in the sociology of knowledge. But the ideal of objectivity need not be thought of as value-free or disinterested. This lecture examines the values, motives, and incentives that animate science and scientists. To what extent are these values cognitive and to what extent is it a problem if they're not? Might the social structure of science generate objective results even if individual scientists are motivated by the pursuit of recognition, money, or tenure? In what ways might the social organization of science be changed in order to increase objectivity? Who should get to participate in the formation of a scientific "consensus" and why? To what extent can the need for scientific expertise be reconciled with the democratic ideal of citizen involvement in important decisions?

Outline

- I. Social factors—money, prestige, political and economic interests, and so on—have often loomed large in the actual practice of science. It has often been implicitly assumed that these social aspects compete with norms of rationality and objectivity that also figure in scientific conduct.
 - A. For the positivists, social factors tend to distort the objectivity that would otherwise result from the application of the scientific method (at least within the context of justification).
 - B. For many of the sociologists of science, appeals to evidence and logic mask the operation of non-evidential interests and biases that constitute the real explanation of scientific conduct.
 - C. We have seen a position between these two views in the work of Kuhn, for whom social aspects of the organization of science can aid, rather than impede, the rationality of science.
- II. Recent work in naturalized epistemology and philosophy of science has followed Kuhn in developing a position according to which social and epistemic norms can cooperate, rather than compete. It has followed the sociologists of science in thinking that even normal science is significantly governed by nonepistemic factors, and it has followed the logical positivists and others in thinking that science is, for the most part, epistemically special.
 - A. It is clear that it can be disastrous for science to be driven by ideology, but it is not clear that ideology need be epistemically harmful to science.
 1. Suppose a classic Marxist critique of science to be entirely correct: Science serves the interests of industrial capitalism. It is plausible that such ideology-driven science would be highly reliable, because industrial capitalism values accurate information about the empirical world.
 2. Such science could count as objective without being disinterested.
 3. This kind of "invisible hand" defense of scientific objectivity will be subject to very severe restrictions, and it does not show that ideology won't lead to scientific distortion. But it's worth noting that ideology doesn't automatically lead to such distortion.
 - B. It can also be argued that the reward structure of science, on the whole, has epistemically salutary effects.
 1. Scientists are rewarded (with prestige, among other things) for having their ideas cited and used. This encourages finding original results and making one's ideas available to others.
 2. Because scientists rely on the ideas of other scientists, the reward system creates some pressure toward testing and replicating the results of others. Ideas are tested through a kind of cooperation and through a kind of competition.
 3. The reward system has some tendency to promote a healthy distribution of scientific labor. If many people are pursuing the most developed research project, it can be rational for other scientists to pursue alternatives.
 4. Although the ordinary self-interest of individuals can lead to a community that functions in a more or less disinterested, inquiring manner, the increasing role of money in science and the recent upsurge in corporate sponsorship of research complicate this model considerably.

- III. Ideology and other sources of idiosyncrasy certainly have exerted embarrassing influence on science in the past, and this observation raises important issues about how objectivity can be cultivated and increased in science.
- A. Individual scientists sometimes evaluate hypotheses partially on the basis of non-evidential factors.
 - 1. Some such factors are politically significant (such as gender, class, race, or nationality), while others arguably stem from considerations as apolitical as birth order.
 - 2. One might expect these non-evidential factors to figure differently in some sciences than in others. Assumptions about gender seem to have crept into primatology but don't seem like much of a worry in theoretical physics. An individual scientist's aesthetic sense might loom large in theoretical physics, however.
 - B. Such protection from distortion and idiosyncrasy as science possesses rests less on finding impartial judges than on structures that bring a range of relevant critical perspectives to bear on ideas and their applications.
 - C. This raises questions about the diversity, in terms of gender, age, birth order, politics, style of intellectual training, and so on, of a given field. Ideally, it seems, you would want as much variety as you could get in order to bring effective criticism in the field.
 - D. The objectivity of a given scientific field is increased by its openness to criticism. Does the field have good conferences and journals? A number of mechanisms can operate to prevent criticism from being as effective as it might be.
 - E. But a version of the "white noise" problem looms large here. Diversity of background and opinion has costs as well as benefits. Requiring evolutionary biologists to take creation scientists seriously might have some tendency to increase the objectivity of the discipline, but it's not clear that it's worth the opportunity costs of doing so.
- IV. Questions about values and the social structure of science loom even larger when we turn our attention to science's role in society at large.
- A. Privately funded science would seem legitimately to serve narrower interests than publicly funded science, but it figures in the public sector when it makes a claim to guide policy or to reveal the truth about something.
 - B. To what extent do scientists have an obligation to reflect on the likely uses of their research? Can one make the argument that the pursuit of knowledge is justified in itself and that the moral consequences should be left to those who apply the research? Much turns on the extent to which benefits and harms of a research project are reasonably foreseeable.
 - C. Issues also arise about how scientists obtain their data. In the United States, if people participate in a medical study, they are owed the highest standard of care. Is it permissible to run studies in other countries for the purpose of avoiding this expensive burden?
 - 1. On the one hand, the researchers seem to be using people as guinea pigs, taking advantage of already significant inequalities.
 - 2. On the other hand, they might well be offering their research subjects better medical care than they would otherwise get. We leave these sorts of issues to ethicists.
 - D. Finally, we note difficulties about scientific decision-making. Nonscientists must rely on scientists to ascertain the scientific significance of such a proposal as the superconducting supercollider. Who should decide whether a supercollider gets built, and how should such decisions be made?

Essential Reading:

Railton, "Marx and the Objectivity of Science," in Boyd, Gasper, and Trout, *The Philosophy of Science*, pp. 763–773.

Longino, "Values and Objectivity," in Curd and Cover, *Philosophy of Science: The Central Issues*, pp. 170–191.

Supplementary Reading:

Godfrey-Smith, *Theory and Reality: An Introduction to the Philosophy of Science*, chapter 11.

Kitcher, *Science, Truth, and Democracy*.

Questions to Consider:

1. Which parts of science seem most and least ideologically driven to you? In the relatively ideological parts of science, to what extent does the presence of ideology undermine objectivity?
2. Which sciences seem to you to strike the best balance between Popperian openness to criticism and Kuhnian consensus about standards and procedures?

Lecture Thirty

Probability

Scope: Through much of Western intellectual history, “chance” was thought to represent the enemy, or at least the limitations, of reason. But notions of chance are now arguably inquiry’s greatest ally. After a potted history of probability, we try to get clear about the basic mathematics of probability, and then we confront the philosophical issues that arise about the interpretation of probability statements. Such statements can be understood in terms of states in the world (for example, relative frequencies) or in terms of degrees of belief (for example, how likely you think it is that the Red Sox will win the World Series).

Outline

- I. Probability has a fascinating history.
 - A. The basic mathematical theory of probability did not really arise until around 1660.
 1. This seems quite shocking, given humanity’s longstanding interest in gambling.
 2. Part of the reason seems to have been that chance did not seem like the sort of thing about which one could have a theory. The traditional Western conception of knowledge as modeled on geometry and as concerning that which must be the case probably played a role. Also, the Christian notion that everything that happens is a manifestation of God’s will may have been a factor.
 3. The study of probability really got going when a nobleman and gambler asked Blaise Pascal to solve some problems about how gambling stakes could be divided up fairly.
 4. Probability caught on very quickly, if somewhat haphazardly, in business, law, and other applications.
 - B. Arguably, probability is crucial to the modern conception of evidence.
 1. The term *probability* started off being associated with testimony. An opinion was probable if grounded in reputable authorities. On that basis, it was not uncommon to hear it said that an opinion was probable but false, meaning that the authorities were wrong in that case.
 2. Probability eventually morphed sufficiently to allow the necessitating “causes” of high sciences, such as physics and astronomy, to be assimilated to the mere “signs” of low sciences, such as medicine. The low sciences, lacking demonstrations, relied on testimony.
 3. It was only in the Renaissance that the notion of diagnosis was distinguished from such notions as authority and testimony, on one hand, and from direct dissections and deductive proof, on the other. Probability becomes evidence when it becomes the testimony of the world, as it were. A symptom testifies to the presence of disease.
 4. As the idea that physics, for example, could be demonstrative like geometry fades, we are left with an idea of evidence that derives from signs and symptoms. We have evidence of when one bit of the world indicates what another bit of the world is like.
 - C. In the 19th century, the spread of probabilistic and statistical thinking gradually undermined the assumption that the world was deterministic.
 1. As governments kept better records of births, deaths, crimes, and so on, it emerged that general patterns could be predicted in a way that individual events could not.
 2. As statistical laws became more useful, the assumption that they reflected underlying but virtually unknowable deterministic laws became increasingly irrelevant. The statistical laws started to seem the stuff of science, not a substitute for real science.
 3. Important parts of statistical thinking migrated from such disciplines as sociology to such disciplines as physics, where, again, supposed deterministic explanations started to seem irrelevant.
 4. With the arrival of quantum mechanics in the early 20th century, we encounter powerful arguments to the effect that our world is governed by statistical laws that are not backed by deterministic ones.
- II. The mathematics of probability is uncontroversial. A somewhat casual sense of the mathematics will be adequate for our purposes.
 - A. All probabilities are between 0 and 1.
 - B. Any necessary truth gets assigned a probability of 1.

- C. If A and B are mutually exclusive, then the probability that one or the other will happen is equal to the sum of their individual probabilities.
 1. If there is a 30% chance that you will have the ranch dressing and a 40% chance that you will have the vinaigrette, then there is a 70% chance that you will have either the ranch or the vinaigrette (assuming, I hope correctly, that you'd never mix the two).
 2. Things get more complicated if the outcomes are not mutually exclusive. The chance that I will have the cake or the pie (given that I might have both) is the chance that I will have one plus the chance that I will have the other minus the chance that I will have both (essentially that is to avoid double-counting).
- D. Other rules for calculating probabilities can be built (roughly) from these.

III. Controversy arises in the interpretation of the mathematics. We will consider three major interpretations.

- A. *Frequency theories* place probabilities “out there” in the world. This is the most commonly used concept of probability in statistical contexts. The frequency theory identifies probabilities with certain relative frequencies.
 1. Probabilities could be construed as *actual* relative frequencies. The probability of getting lung cancer if you smoke is the ratio of smokers with lung cancer to the total number of smokers. This approach is clear and links probabilities tightly to the evidence for them.
 2. This approach faces issues about how to place objects in scientifically salient populations. The probability that I will get lung cancer is either 1 or 0. And I have one probability of getting lung cancer as a nonsmoker, another as a 40-year-old male, another as a coffee addict, and so on.
 3. A more serious problem occurs because this account is “too empiricist.” It links a scientific result too closely to experience. A coin that has been tossed an odd number of times cannot, on this view, have a probability of .5 of coming up heads. In addition, a coin that has been tossed once and landed on heads has, on this view, a probability of 1 of landing on heads. Such single-case probabilities are a real problem for many conceptions of probability.
 4. One might go with *hypothetical* limit frequencies: The probability of rolling a seven using two standard dice is the relative frequency that would be found if the dice were rolled forever. We saw an idea like this in the pragmatic vindication of induction.
 5. This version might not be empiricist enough. The empiricist will want to know how our experience in the actual world tells us about worlds in which, for example, dice are rolled forever without wearing out.
- B. *Logical theories* treat probabilities as statements of evidential relationships. They can be interpreted as the judgments of an ideal agent or as relations in logical space. The idea here is that probability gives a logic of partial belief or inconclusive evidence modeled on what deductive logic provides for full belief or conclusive evidence.
 1. Just as our full beliefs should not contradict one another, our partial beliefs should cohere with one another. Having coherent beliefs is not sufficient for getting the world right, but having incoherent beliefs is sufficient for having gotten part of it wrong.
 2. Probabilistic coherence is a matter of how well an agent's partial beliefs hang together. If your evidence assigns a probability of .8 to p , then it had better assign a probability of .2 to *not-p*.
 3. Logical theories of probability impose conditions beyond mere coherence. In particular, they impose the *principle of indifference*. If your evidence does not give you a reason to prefer one outcome to another, you should regard them as equally probable.
 4. The mathematics of probability does not require this principle, and it turns out to be very troublesome. There are many possible ways of distributing indifference, and it's hard to see that rationality requires favoring one of these ways.
- C. *Subjective theories* treat probabilities as degrees of belief of actual agents—they directly concern the believing agent rather than the world, but they are subject to objective although rather minimal criteria of rationality.
 1. A degree of belief is measured by one's notion of a fair bet. The odds at which you think that it would be reasonable to bet that a Democrat will win the next presidential election tell you the extent to which you believe that a Democrat will win.

2. Because this approach does not explicate probabilities in terms of frequencies or a principle of indifference, it relies only on the notion of probabilistic coherence to make probability assignments “correct.”
3. For this reason, this model as so far described seems to allow any old probabilistically coherent set of beliefs to be perfectly rational. Paranoid delusions tend to be strikingly coherent yet seem to be rationally criticizable. We will address this problem in the next lecture.

Essential Reading:

Curd and Cover, “Bayes for Beginners,” in Curd and Cover, *Philosophy of Science: The Central Issues*, pp. 627–638.

Hacking, *The Emergence of Probability*.

Supplementary Reading:

Hacking, *The Taming of Chance*.

Questions to Consider:

1. Geometry served as a paradigm of knowledge for centuries. What paradigms of knowledge operate in our culture at present? Do any of them reflect the shift discussed in this lecture to the idea that we can have knowledge of contingent matters?
2. If you knew that an urn consisted of red and green balls (but knew nothing else about it), would it be irrational to let the fact that you like red better than green affect your probability judgments? What kind of mistake, if any, would you be making if you assigned a probability of .9 to drawing a red ball and .1 to drawing a green one?

Lecture Thirty-One

Bayesianism

Scope: Bayesian conceptions of probabilistic reasoning have exploded onto the philosophical and scientific scene in recent decades. Such accounts combine a subjectivist interpretation of probability statements with the demand that rational agents update their degrees of belief in accordance with *Bayes's Theorem* (which is itself an uncontroversial mathematical result). Bayesianism is a remarkable program that promises to combine the positivists' demand for rules governing rational theory choice with a Kuhnian role for values and subjectivity. After explaining and motivating the basics of Bayesianism, we examine its approach to scientific theory choice and to the raven paradox and the new riddle of induction.

Outline

- I. Starting from very modest resources, the *Bayesian* approach to probability has rejuvenated philosophical thinking about confirmation and evidence.
 - A. Bayesianism begins with a subjective interpretation of probability statements: They characterize personal degrees of belief. These degrees of belief can be more or less measured by betting behavior; the more unlikely you think a statement is, the higher the payoff you would insist on for a bet on the truth of the statement.
 - B. Your degrees of belief need not align with any particular relative frequencies, and they need not obey any principle of indifference. Bayesianism requires little more than probabilistic coherence of beliefs.
 - C. The *Dutch book argument* is designed to show the importance of probabilistic coherence. To say that a Dutch book can be made against you is to say that, if you put your degrees of belief into practice, you could be turned into a money pump.
 1. If I assign a .6 probability to the proposition that it will rain today and a .6 probability to the proposition that it will not rain today, I do not straightforwardly contradict myself.
 2. The problem emerges when I realize that I should be willing to pay \$6 for a bet that pays \$10 if it rains, and I should be willing to pay \$6 for a bet that pays \$10 if it does not rain.
 3. At the end of the day, whether it rains or not, I will have spent \$12 and gotten back only \$10. It seems like a failing of rationality if acting on my beliefs would cause me to lose money no matter how the world goes.
 4. It can be shown that if your degrees of belief obey the probability calculus, no Dutch book can be made against you.
- II. But pretty loony webs of belief can still be probabilistically coherent. Bayesianism becomes a serious scientific theory of scientific rationality by developing a theory of how one should handle evidence. The first component of this theory is a notion of confirmation as raising the probability of a hypothesis.
 - A. Bayesians think that the notion of confirmation is inherently quantitative. We cannot ask whether a piece of evidence, E, confirms a hypothesis, H, unless we know how probable H started out being—we have to have a *prior probability* for H. E confirms H just in case E raises the prior probability of H. This means that the probability of H given E is higher than the probability of H had been: $P(H/E) > P(H)$. E disconfirms H if $P(H/E) < P(H)$.
 - B. All this is done within the subjectivist or personal interpretation of probability. A big cloud on an otherwise clear horizon counts as evidence of rain for me, just in case my subjective probability that it will rain, given the new information that there is a big cloud on the horizon, is higher than my prior probability that it would rain.
 - C. In saying this, we have made tacit use of the notion of *conditional probability*: the probability of the hypothesis conditional on or given the evidence.
 1. The conditional probability of H given E is the probability of (H&E) divided by the probability of E (provided that E has a nonzero probability). (H&E) is the intersection, the overlap, of cloudy days and rainy days. The definition says that the higher the percentage of cloudy days that are rainy, the higher the conditional probability of H given E.

2. If I were already convinced that it would rain (because of a weather report, for instance), then this high conditional probability of rain depending on clouds would not change my prior belief and, thus, would not be evidence. But, if I had been relatively neutral, it might significantly confirm the rain hypothesis for me.
- D. The idea that whatever raises the probability of H confirms H is not without its problems. My seeing Robert De Niro on the street might raise the probability that he and I will make a movie together, but it hardly seems to count as evidence that we'll make that movie. However, we'll assume that such problems can be solved.
- III. The second idea crucial to Bayesians is that beliefs should be updated in accordance with *Bayes's Theorem*.
- A. The theorem itself is a straightforward consequence of the definition of conditional probability. Non-Bayesians accept the truth of the theorem but don't put it to the use that Bayesians do.
- B. The classic statement of the theorem is: $P(H/E) = \frac{P(E/H) \times P(H)}{P(E)}$.
- C. The left side of the statement is the conditional probability of the hypothesis given the evidence. It can have two different readings, depending on whether the evidence is "in" yet or not.
1. If the evidence is not in, then $P(H/E)$ is the prior conditional probability of H given E. If I were a physicist in 1915, I might have assigned a low probability to Einstein's hypothesis of general relativity, but I also might have thought to myself, "If it turns out that light rays are bent by the Sun, I assign a quite high probability to Einstein's hypothesis."
 2. If the evidence is in, then $P(H/E)$ represents the posterior probability of the hypothesis. It is the probability I now assign to Einstein's hypothesis, once I have gotten news that light rays are bent.
- D. We now unpack the right side of the statement.
1. $P(E/H)$ measures how unsurprising the evidence is given the hypothesis. Given Einstein's hypothesis of general relativity, the probability that light rays are bent by the Sun's gravitational field is quite high.
 2. $P(H)$ is just the prior probability of the hypothesis.
 3. The posterior probability (that is, the left side of the equation) is directly proportional to the prior probability of the hypothesis and directly proportional to the extent to which the hypothesis makes evidence unsurprising.
 4. The prior probability of the evidence is the denominator of the fraction, reflecting the fact that, all other things being equal, unexpected evidence raises posterior probabilities a lot more than expected evidence does. Apart from Einstein's theory, the probability of light being bent by the Sun was quite low. It is because Einstein's prediction is so unexpected, except in light of Einstein's theory, that the evidence had so much power to confirm the theory.
 5. Thus, the more unexpected a given bit of evidence is apart from a given hypothesis and the more expected it is according to the hypothesis, the more confirmation the evidence confers on the hypothesis.
- E. The controversial part arises when the Bayesian proposes as a rule of rationality that, once the evidence comes in, the agent's posterior probability for H given E should equal the agent's prior conditional probability for H given E.
1. This sounds uncontroversial; as we saw, there were just two interpretations of the left side of the equation. But the mathematics by itself will not get you this result.
 2. Once the evidence comes in, I could maintain probabilistic coherence by altering some of my other subjective probabilities, namely, some of the numbers on the right side. I could decide that the evidence was not that surprising after all, for instance, thereby making my posterior probability different from my prior conditional probability. Why must today's priors be tomorrow's posteriors?
- F. The Bayesian appeals to a diachronic (across time) Dutch book argument to support this requirement. If you use any rule other than Bayesian conditionalization to update your beliefs, then a bookie who knows your method can use it against you by offering you a series of bets, some of which depend on your future degrees of belief.

- IV. Bayesianism has helped rekindle interest in issues about evidence and justification. The Bayesian approach allows for impressive subjectivity (there are very few constraints on prior probabilities other than coherence with other degrees of belief) and impressive objectivity (there is one correct way of updating one's beliefs in the face of new evidence).
- A. Bayesians argue that initial subjectivity disappears when enough good evidence comes in. This is called the *washing out of prior probabilities*. It can be established that no matter how great the disagreement is between two people, there is some amount of evidence that will bring their posterior probabilities as close together as you would like. That is impressive, but it is subject to some significant limitations.
1. If one person assigns a prior probability of 0 to a hypothesis, no evidence will ever increase that probability.
 2. There is no assurance that convergence will happen in a reasonable amount of time.
 3. The washing-out results require that the agents agree about the probabilities of all the various pieces of evidence given the hypothesis in question. This seems problematic.
- B. Bayesianism's attractiveness as a theory of scientific inference can be appreciated by revisiting Goodman's new riddle of induction and Hempel's raven paradox.
1. The Bayesian will say that there is nothing the matter with either of the new riddle's inductive arguments. It is fine to infer from the greenness of emeralds to their continued greenness or from their "grueness" to their continued "grueness." Whichever hypothesis you think more probable going in will remain more probable going out.
 2. Bayesians can handle the raven paradox equally straightforwardly. The greater the ratio of $P(E/H)$ to $P(E)$, the greater the power of evidence to confirm H . This turns out to be the source of the difference in the confirming power of white shirts and black ravens to confirm "All ravens are black." The probability that the next raven I see will be black given that all ravens are black is 1. The probability that the next shirt I see will be white given that all ravens are black is much lower. It is pretty much just my prior probability that the next shirt I see will be white.

Essential Reading:

Salmon, "Bayes's Theorem and the History of Science," in Balashov and Rosenberg, *Philosophy of Science: Contemporary Readings*, pp. 385–404.

Godfrey-Smith, *Theory and Reality: An Introduction to the Philosophy of Science*, chapter 14.

Supplementary Reading:

Salmon, "Rationality and Objectivity in Science or Tom Kuhn meets Tom Bayes," in Curd and Cover, *Philosophy of Science: The Central Issues*, pp. 551–583.

Questions to Consider:

1. Utter fictions can be quite coherent. Does the Bayesian need an argument that a set of beliefs that is probabilistically coherent (both at a time and across time) is likely to be true? Does the Bayesian have resources to provide such an argument?
2. Does the Bayesian solution to Goodman's new riddle of induction seem satisfactory to you? The solution works if you have the right prior probabilities, but it doesn't claim that you *should* have those prior probabilities.

Lecture Thirty-Two

Problems with Bayesianism

Scope: Predictably, a Bayesian backlash has also been gaining momentum in recent years. This lecture investigates Bayesianism's surprisingly subjective approach to probability assignments, as well as the Bayesian treatment of the problem of old evidence (it appears that we can never learn anything from evidence that is already in). We compare the Bayesian approach with competing conceptions of statistical inference, such as those derived from classical statistics. This assessment results in a cost-benefit analysis rather than a vindication or a refutation.

Outline

- I. Though mathematically intensive, the basic ideas behind Bayesianism are rather simple and powerful. It has gained many adherents in recent decades, but with increased attention has come increased criticism. We begin with a couple of criticisms that we will not pursue in detail.
 - A. Bayesianism involves a rather dramatic idealization of human cognizers.
 1. We do not have the processing power to meet Bayesian standards even in fairly simple cases. Coherence requires logical omniscience, namely, that we know all the logical consequences of our beliefs, and that is unrealistic.
 2. On the other hand, it is not clear how descriptively accurate the theory needs to be. There is at least some role for ideals that cannot be met.
 - B. Many think that Bayesianism does not reflect actual scientific practice. Scientists do not think of their work in terms of degrees of belief. They leave themselves out of the picture when doing science.
- II. The *problem of old evidence* represents a longstanding challenge to the Bayesian approach.
 - A. It seems that scientific theories can be confirmed by facts that are already known. Newton's theory, for instance, could explain Kepler's well-known laws, and thus, Kepler's laws are evidence for Newton's theory.
 - B. But any evidence that is already known for sure should, it seems, receive a probability of 1. And because $P(E)$ is 1, $P(E/H)$ is 1.
 - C. If we plug these numbers into Bayes's Theorem, we quickly see that old evidence has no power to confirm hypotheses. The prior probability of the hypothesis is multiplied by 1 over 1, so the posterior probability stays equal to the prior probability.
 - D. A couple of responses are open to the Bayesian.
 1. Bayesians can claim that the subjective probability of E should not be considered against one's actual background knowledge (because that knowledge includes E) but, instead, against what the background knowledge would be if E were not yet known. This involves ascertaining how surprising E would be if neither it nor anything that entails it were included in our background knowledge.
 2. But the counterfactual "how surprising would I judge E to be if I did not already know it" can be difficult to evaluate once we realize how many statements might bear on E.
 3. Alternatively, the Bayesian can say it's not really E that confirms H when E is already known. In the Newtonian case, it was the new information that Newton's theory entailed Kepler's laws that did the confirming. It is "H entails E," not E, that confirms.
 4. This involves a couple of problems: Sometimes, the fact that H entails E itself seems like old evidence. And anyway, mightn't we want to insist that E confirms H in such circumstances?

- III. The most influential objections to the Bayesian program concern its somewhat brazen tolerance of subjective probabilities. In scientific contexts, Bayesianism can be supplemented with various strategies or scientific values designed to impose constraints on admissible probability of conditional probability assignments so that outlandish degrees of belief are regarded as legitimately criticizable. The extent to which subjectivity can be tempered is different for the various terms of Bayes's Theorem.
- A. $P(E/H)$ is usually pretty well behaved. Often, the hypothesis in question entails the evidence, in which case $P(E/H)$ is 1. The probability of this piece of copper conducting electricity given that all copper conducts electricity is 1.
 - B. The prior probability of the hypothesis can be tamed a bit.
 1. One might try to impose norms requiring one to look for evidence (for example, observed frequencies) relevant to setting prior probabilities, and one might evaluate new hypotheses by comparing them to similar hypotheses.
 2. This is trickier than it sounds. What is to count as a hypothesis that is similar to the one in question? Evidence seems unlikely to settle the relevant similarity relation.
 - C. The hardest problem concerns the denominator of Bayes's Theorem— $P(E)$.
 1. The probability of the evidence is equal to the probability of the evidence on the assumption that our hypothesis is true plus the probability of the evidence on the assumption that our hypothesis is false: $P(E) = [P(E/H) \times P(H)] + [P(E/\sim H) \times P(\sim H)]$.
 2. The claim that our hypothesis is false is not itself a hypothesis. It is called the *catch-all hypothesis*. There are endless ways in which our hypothesis could be false, and they will not all assign the same probability to the evidence.
 3. The only way to get solid evidence for the values in this part of the equation is to claim that all the possible hypotheses are under consideration. And this is generally not warranted.
 4. Thus, even if you can temper or constrain many of the probabilities that figure in the theorem in the light of evidence and scientific practice, there is no getting around the fact that one of the probabilities in the equation can be only a kind of guess about how surprising a certain result would be. Subjectivity of this sort cannot be eliminated if you're going to use Bayes's Theorem.
 - D. Though many scientists think that nothing having anything to do with subjectivity should be let anywhere near science, it's not clear how bad the subjectivity built into the Bayesian program is. Kuhn, for instance, thought a limited role for subjectivity was crucial to the health of science.
- IV. Classical statistics tries to avoid this role for subjectivity. It remains dominant in most scientific disciplines, though Bayesianism is on the rise.
- A. If you don't start from probabilities for hypotheses, you can't end up with them. Because non-Bayesians don't like the role of subjectivity in setting $P(H)$, they forego getting any values for $P(H/E)$.
 - B. Basically, that leaves them with $P(E/H)$ doing most of the work.
 1. In classical statistics, we want to know whether a given correlation is significant or random.
 2. We start by assuming that the results are random, and we run some tests.
 3. Very roughly, if $P(E/H)$, that is, the probability of getting results like these randomly is lower than .05, then we reject the assumption of randomness and call the correlation significant.
 4. This significance threshold is both somewhat arbitrary and somewhat sacred in science. One might ask exactly why rules that are arbitrary but objective are better than probability judgments that are subjective but objectively updated; however, we can only touch on these issues here.
 5. The 5% threshold has occasionally been abused or applied mindlessly, and that has led to some very questionable scientific work.
 - C. Another statistical confusion looms large in public discussions of evidence. All sides agree that it is important to keep $P(E/H)$ quite distinct from $P(H/E)$, but people make this mistake quite often.
 1. In a criminal trial, the jury might be told an impressive $P(E/H)$, for instance, that the forensic evidence matches the defendant and that the chances of its matching a person chosen at random are miniscule.
 2. This evidence has a lot of potential power to confirm the hypothesis that the defendant is guilty, but it cannot do so without a prior probability for that hypothesis. If I were having dinner with the defendant 1,000 miles away from the scene of the crime, the evidence will rightly fail to sway me. It's only by assuming certain values for the prior probability of H that this argument goes through.

3. Clearheaded approaches in the tradition of classical statistics understand that one can never get directly from $P(E/H)$ to any value for $P(H)$. Recent developments in the field try to allow the classical approach to do a lot more than apply mechanical rules of rejection, while still avoiding a role for prior subjective probabilities. But they cannot provide posterior probabilities, which may be a bug or a feature.

Essential Reading:

Glymour, “Why I Am Not a Bayesian,” in Curd and Cover, *Philosophy of Science: The Central Issues*, pp. 584–606.

Supplementary Reading:

Kelly and Glymour, “Why Probability Does Not Capture the Logic of Scientific Justification,” in Hitchcock, *Contemporary Debates in Philosophy of Science*, pp. 94–114.

Questions to Consider:

1. You and I are not capable of meeting the demands of Bayesian rationality—we simply don’t possess adequate computing power. Is this an objection to Bayesianism as a theory of scientific reasoning or not? Is it an objection to a moral theory if you and I aren’t capable of living up to its demands? Why or why not?
2. What would be the pros and cons of instructing jurors to think of themselves as updating prior subjective probabilities on the basis of the evidence presented to them?

Lecture Thirty-Three

Entropy and Explanation

Scope: Most philosophy of science these days is philosophy of a particular science and, more particularly, of a particular issue or theory within one of the sciences. As we wind down the course, I will try to offer some illustrations of how the general issues in philosophy of science that we have discussed are being treated within contemporary, relatively specialized philosophy of science. In this lecture, we turn to the philosophy of physics and examine an intriguing package that includes the reduction of thermodynamics to statistical mechanics, the direction of time, the origin of the universe, and the nature of explanation.

Outline

- I. We turn now to a series of relatively detailed examinations of philosophical issues that arise within particular sciences. These both illuminate and are illuminated by the general philosophical issues on which we've so far focused.
 - A. Our topic from the philosophy of physics can seem frivolous (especially the way I've chosen to express it), but it raises deep issues about explanation and reduction. Why can I not stir milk out of my coffee?
 - B. In one sense, I *can* stir milk out of my coffee. The basic laws of nature (both classical and quantum mechanical) permit it. They permit all of the gas in a container to cluster in one corner, and they permit heat to flow from a metal bar that has been kept in the freezer to one that has been kept in a hot oven.
 - C. There is nothing in the basic laws of motion specifying in which direction molecules must move—a reverse motion is permitted by the basic mechanics.
 - D. But, though they are permitted by the basic laws—in this case, statistical mechanics—the laws that in some sense reduce to these more basic laws tell us that air never leaks into a punctured tire.
 - E. The second law of thermodynamics says that energy tends to spread out. Another way of saying this is that *entropy* (a measure of this dissipative tendency) tends to increase. A system that is energetically isolated (energy is neither added to nor removed from the system) will tend to move toward an equilibrium state—heat will spread out and stay spread out.
- II. Why does the second law of thermodynamics hold, given that there is nothing about the laws of motion, taken just by themselves, to make it so? The 19th-century Austrian physicist and mathematician Ludwig Boltzmann worked out the two most influential answers to this question.
 - A. Boltzmann's first answer is that the effect of collisions between rapidly moving gas molecules would tend to bring about an increase in entropy until it reached its maximum value. As Boltzmann realized, his answer was statistical, not deterministic.
 - B. Boltzmann's second answer suggests that there are just more ways for particles to spread out than there are for them to be concentrated. This can be thought of as a matter of multiple realizability; high-entropy states are realized by many more lower-level states than low-entropy states are. This explanation also makes the second law statistical.
 - C. The first explanation provides a mechanism for the tendency for entropy to increase; the collisions bring about the increases in entropy.
 - D. The second approach explains without providing a mechanism. Just as you don't need a causal story about the shuffling of cards to understand why you never get dealt a royal flush, the causal details are largely irrelevant to the second explanation.
- III. Whichever explanation we adopt, we will run into puzzles about the direction of time, however.
 - A. Let's take the second explanation first. Just as almost all of the states a closed system can move to are high-entropy states, just about all the states it could have moved from are high-entropy states. Thus, it would seem that entropy should increase as we move toward the past, just as it does when we move toward the future. But that never happens. There is a *temporal asymmetry* at the observable, thermodynamic level that this explanation does not seem to account for.

- B. The case with the first, more mechanical explanation is a bit more complicated.
 1. As we've seen, the basic laws of motion permit the collisions to "run backward."
 2. Still, if we can appeal to facts about collisions to provide a mechanism for entropy to increase, then we finally have a time asymmetry built into our system. Oddly, no mechanical account of how entropy is brought about has, to my knowledge, gained widespread assent.
 3. Thus, it's not clear that we have a mechanical explanation of the direction of time, and the more purely statistical explanation, as we've seen, would lead us to expect entropy to increase in both temporal directions.
- C. The problem, then, is not merely that the laws of motion say that decreasing entropy should be possible but that it does not happen. Possible things fail to happen all the time. The problem is that thermodynamics, in particular the second law, seems to be in conflict with what the underlying laws of statistical mechanics would lead us to expect.
- D. Here's a similar way of approaching the problem: If states of thermodynamic equilibrium are overwhelmingly the most probable ones, why is the world we observe so full of situations that are so far from equilibrium?
 1. Boltzmann suggests that we inhabit a peculiar corner of the universe where the thermodynamic equilibrium states that hold sway in most parts of the universe do not obtain.
 2. Only very peculiar combinations of circumstances will give rise to organisms that can think and observe. The reason we always see entropy increasing is because we inhabit a corner of the universe in which entropy is abnormally low. It has got nowhere to go but up.
 3. In other parts of the universe, entropy might decrease as much as it increases. Boltzmann suggests that this raises deep questions about the direction of time in those parts of the universe.
- IV. The most influential answer to the puzzle about why entropy seems always to increase generalizes Boltzmann's suggestion beyond our corner of the universe. Entropy is on the rise everywhere because it started out very low everywhere.
 - A. Even the mechanical explanation for the tendency of entropy to increase needs to posit a low-entropy state in the past. If entropy had started high, the mechanisms would help keep it there, but they wouldn't account for the overwhelming tendency for entropy to increase that we seem to observe.
 - B. Thus, it seems that we must adopt the *Past Hypothesis*, according to which entropy started very, very (add about 10 to the 23rd power *very's*) low. If the universe is constantly moving toward more probable states, it must be moving from a *mighty* improbable state.
- V. At this point, a major issue in the philosophy of explanation arises: Does the Past Hypothesis need to be explained?
 - A. Here's a way of fleshing out the Past Hypothesis that makes it seem to cry out for explanation.
 1. Matter seems weirdly uniformly distributed around 100,000 years after the Big Bang. When dealing with an attractive force such as gravity, a uniform distribution of matter is highly unusual, because objects will tend to clump together.
 2. Huw Price, a philosopher of science, compares the Past Hypothesis to the idea of throwing trillions of foam pellets into a tornado and having them shake down into a uniform sheet, one pellet thick, over every square centimeter of Kansas. The Past Hypothesis differs from this mainly in being enormously less probable (according to some calculations, anyway).
 3. Furthermore, says Price, the Past Hypothesis is the only weird initial condition that we need in order to account for all of the low-entropy systems in the universe, because the improbable initial smoothness in the universe led to the formation of stars and galaxies, and these sorts of things are responsible for the temporarily asymmetric phenomena we encounter.
 4. Thus, deep facts about our universe seem to turn on an enormously improbable fact, namely, the incredibly low-entropy state of the universe at a certain point relatively soon after the Big Bang. Surely something that important and that improbable needs to be explained.
 - B. But there are powerful reasons for wondering what could possibly explain such a fact and for wondering whether such an explanation would ultimately be scientific. Some serious empiricist worries loom large here.

1. The Past Hypothesis can be compared to the *First Cause* argument for God's existence, and similar worries arise. Why carry the demand for explanation this far and no further?
 2. It will not help to explain a past state of surprisingly uniform distribution of matter by positing an even more improbable state before that.
 3. There's also a worry about initial conditions and single-case probabilities here. If universes were as plentiful as blackberries, we could pursue explanatory hypotheses about how they arise and develop.
- C. Does the Past Hypothesis count as a law? It is a prime example of something that happens only once that might still count as a law.
1. It does not have the logical form we associate with laws. But it functions crucially in explanation of many different phenomena; thus, it might count as a law on a "systems" conception of laws, which identifies laws with axioms of true deductive systems that best combine strength and simplicity.
 2. Some think that calling the Past Hypothesis a law makes it stand in less need of explanation.

Essential Reading:

Price, "On the Origins of the Arrow of Time: Why There Is Still a Puzzle about the Low-Entropy Past," in Hitchcock, *Contemporary Debates in Philosophy of Science*, pp. 219–239.

Callendar, "There Is No Puzzle about the Low-Entropy Past," in Hitchcock, *Contemporary Debates in Philosophy of Science*, pp. 240–255.

Supplementary Reading:

Sklar, *Philosophy of Physics*, chapter 3.

Questions to Consider:

1. Does everything that is highly improbable call out for explanation? Why or why not? And which sense of probability (frequency or degree of belief, for example) figures in the notion of improbability at work here?
2. It is sometimes said that unique events cannot be explained, perhaps because explanation involves placing events in a pattern. But don't we sometimes explain unique occurrences? How do we do so?

Lecture Thirty-Four

Species and Reality

Scope: What kind of a kind is a species, if indeed it is a kind at all? We certainly talk as if species have properties of their own (such as being endangered), and in fact, a species is, in many respects, more like an individual than it is like a class or kind. But biology defines species in a number of ways, and even some of the best definitions seem to exclude most organisms on Earth from being members of a species. In this lecture, we try to understand the motivations behind biological classification, and we wonder about the things so classified. How are we to decide whether a species concept is a good one? And how are we to decide whether a good species concept tracks something real?

Outline

- I. We now turn to the philosophy of biology, probably the most rapidly expanding field within the philosophy of science. We will focus on the notion of a species and use that as a window into parts of the philosophy of biology and into general philosophy of science issues about classification and reality of scientific kinds or classes.
 - A. The species concept figures centrally in biology.
 1. Species are fundamental units in the story of evolution. They are born, split into new species, and eventually become extinct.
 2. They are also fundamental to biological classification. Members of the same species have something biologically important in common.
 3. As uncontroversial as these claims sound, together, they put real pressure on the notion of a species. It is not obvious that the notion can serve both these functions well. The properties of organisms and populations that are relevant to the story of evolution might be different from those that are important for certain classificatory purposes.
 4. Furthermore, species loom large in our applications of biological thinking, for example, in some environmental protection laws.
 - B. The very notion of a species can be made to seem puzzling.
 1. Evolutionary change is, in some important sense, gradual; new kinds of creatures arise via small mutations in existing creatures. Where, then, are we to find distinctions of kind within a fundamentally continuous process?
 2. Just as there is impressive continuity across species, there is impressive variety within species. *Conspecifics* (members of the same species) are not united by a common essence. No genetic, phenotypic, or behavioral trait is essential to making something a member of a species. Nor are there, in general, traits that are unique to a given species.
- II. This raises the issue of the ontological status of species; what kind of entity is a species? One surprising but common answer is that species are more like individuals than like classes of objects.
 - A. A kind of spatiotemporal or causal connectedness is required for a species. If evolutionary processes are primary, it is plausible to hold that something has to be part of a lineage to be part of a species. And a lineage is a particular thing, not an open-ended class of things.
 - B. On this view, species have a beginning and an end in time, and they have a spatial location.
 - C. Perhaps most importantly, species are constituted by properties at the population level, not at the level of the individual organism. This gives a species a certain cohesiveness needed to play a role in scientific explanations. It is the population, not individual organisms, that bears such properties as “having lost much of its genetic diversity” or “having a trait that was rare become prominent.” These are taken to be genuine and explanatorily important biological properties.
 - D. These population-level properties can change rapidly, in accordance with our sense that boundaries between species are relatively stark. This is part of the explanation of how gradualism and continuity at the level of parent-offspring genetic relationships can be reconciled with the idea that organisms seem to come in distinct kinds.

III. A surprising number of definitions of *species* have been proposed by biologists.

- A. *Phenetic species concepts* group organisms in terms of genetic, behavioral, or morphological similarity.
 - 1. But similarity, as we've seen in this course, is a tricky notion. How is it to be measured, for instance?
 - 2. Intraspecific similarity is less impressive than one might think. Queen and drone bees don't appear all that similar.
 - 3. If the notion of a species is determined by similarity, then species membership can't be used to explain similarity.
- B. According to the *biological species concept*, a species is a group of organisms that can interbreed and produce fertile offspring. A species is characterized by the relatively free flow of genes within it.
 - 1. This concept is difficult to apply over time. Let A, B, and C be members of successive generations. Suppose that B is about equally similar to A and to C and could breed with either one but that A and C would not be able to breed. If A is the standard, then C is a member of a new species, but if B is the standard, there is just one species. Proponents of this species concept apply it only at a time, not over time. But that means they need an independent notion of a speciation event in order to determine whether a creature is conspecific with a given ancestor.
 - 2. The notion of a reproductively isolated population is also tricky. Not all kinds of reproductive isolation (such as being kept in a zoo) count.
 - 3. The most glaring problem with the biological species concept is that, as it stands, it does not even apply to creatures that reproduce asexually. Furthermore, gene flow is much easier among plants and single-celled organisms than it is in multicellular animals.
- C. *Phylogenetic species concepts* define species in terms of a shared history. The biological concept involves a theory about the process whereby species are created and sustained, while phylogenetic accounts simply appeal to patterns of common ancestry.
 - 1. They thus make room for the idea that mechanisms other than reproductive isolation can produce speciation.
 - 2. So far, this is just a *grouping criterion*: It lumps organisms together in ways that matter to evolution, but it does not provide *ranking criteria*; it doesn't tell us which groups are species or genera or sub-species and so on.
- D. The *ecological species concept* identifies a species as a group of organisms sharing a particular ecological niche.
 - 1. This looks like an attractive way to handle asexually reproducing species, because they do not compete with one another for mating opportunities, but they do for roles within the ecosystem.
 - 2. But how well do we understand ecological niches and how enduring are they?

IV. How important is it to unify these various conceptions of species?

- A. Monists think there is a single correct species concept, but monism runs the risk of excluding species concepts with genuine explanatory power. Might multiple ways of identifying species each answer to legitimate scientific purposes?
 - B. Pluralists think that there is no problem having multiple conceptions of species. Pluralism comes in degrees, but the more tolerant one is of different species concepts, the more an explanation seems to be needed of what makes each of them a *species* concept.
 - C. Some thinkers are skeptics about species. They deny that anything in the world answers to all the uses to which the notion of species gets put.
- V. These issues about the reality of a grouping or category arise even more clearly in discussions of higher taxa, such as families, groups, and genera. As with species, pluralists will stress the legitimacy of different purposes served by classification, while monists will point out conceptual problems that seem to stand in the way of thinking of different classifications as right or real.
- A. A *phenetic* classification system would try to convey information about similarity in a maximally efficient way. Not only are species maximally similar organisms, but genera are maximally similar species, families are maximally similar genera, and so on.

- B. A related but distinct approach would be to classify in terms of *evolutionary disparity*. This is simultaneously a historical and morphological system. This approach might accord a lizard species sufficiently different from all other species of its own genus.
- C. Finally, we might classify in a way that reveals evolutionary history—organisms classified in terms of ancestry. This is the *cladistic* approach to classification.
- D. Pluralists might be tempted to allow all three kinds of classification.
- E. Monists (and others) might object to the phenetic and evolutionary disparity systems on the grounds of the unclarity of similarity relation. They might also object that the cladistic approach can't make room for groups not united by a common ancestry but nevertheless partaking of a genuine explanatory role in evolution.
- F. Cladism is easily the most popular approach to classification, and it has a theory of which biological groups are real—those that share an ancestral species. A reptile is not a real category for cladists because there is no species that is ancestral to all reptiles that is not also ancestral to birds. But this notion of real groups does accord any status to such levels as genera and families. All groups that share a common ancestor species are legitimate and all that do not are not legitimate.
- G. Cladists, monists, pluralists, skeptics, and others all appeal to implicit notions of what makes a group or a distinction real. This discussion helps flesh out our distinction between hard and soft realisms back in Lecture Twenty-Six.

Essential Reading:

Sterelny and Griffiths, *Sex and Death: An Introduction to Philosophy of Biology*, chapter 9.

Supplementary Reading:

Sober, *Philosophy of Biology*, chapter 6.

Questions to Consider:

1. Would you value a distinctive group of animals less if you were to be convinced that it constituted a sub-population but not a species? How, if at all, do various species concepts hook up to what we *value* about species?
2. To a committed defender of a phylogenetic species concept, there is no such thing as a reptile, because the animals we call reptiles don't share a distinctive common ancestor. Does this convince you that the category "reptiles" is illegitimate? Why or why not?

Lecture Thirty-Five

The Elimination of Persons?

Scope: In most cases of reduction, the entity or theory that gets reduced is still presumed to exist; we don't get rid of water by reducing it to H₂O. But in some cases, an *eliminative reduction* seems to be in order. Our best theory of demonic possession says that it never happened; every case of demonic possession is really a case of something else. A number of philosophers have adopted this attitude toward *folk psychology*, the commonsense explanation of behavior in terms of beliefs, desires, and such. Could arguments from neuroscience and philosophy of science really show that there are no such things as beliefs, desires, and persons?

Outline

- I. Arguably, *folk psychology*, our commonsense approach to psychological phenomena, amounts to an ambitious explanatory theory of human behavior. It has a systematic structure deployed for purposes of prediction and explanation.
 - A. The theory has an ontology: It posits unobservable states, such as beliefs, desires, emotions, and so on.
 1. The theory models thoughts on publicly observable (written or uttered) sentences. Beliefs are similarly structured.
 2. Sensations are not taken to be linguistically structured, as thoughts are. They are posited internal states modeled on external objects.
 - B. Folk psychology has laws.
 1. Some of the laws are fairly closely tied to observation, for example, "People who are angry are easily irritated."
 2. Some laws are relatively distant from interpretation via observables, for example, "People will generally choose the means they believe most effective in realizing their ends."
 - C. The theory can be understood using any number of tools we've developed in this course.
 1. It can be structured as an axiomatic system given an interpretation through an observational vocabulary, à la the received view.
 2. Some Kuhnian exemplary applications could be added, along with some discussion of how we learn to predict, explain, and solve puzzles using this theory.
 3. We sometimes use models, as in the semantic conception of theories, to understand folk psychology.
- II. Though we seem to do pretty well predicting and explaining each other's behavior using folk psychology, many philosophers think that folk psychology competes with and compares unfavorably to various kinds of scientific psychology.
 - A. The laws that figure in folk psychological explanations can be saved from the "death of a thousand counterexamples" only by being protected by lots of "all-other-things-equal" clauses or by being formulated in terms of tendencies people have. They are hardly bold Popperian conjectures, and we tend to explain away their failure by appealing to ad hoc hypotheses.
 - B. Furthermore, folk psychology has not made much progress in the last few thousand years. Competing research programs look more progressive.
 - C. Folk psychology does not explain many phenomena that appear to fall within its domain: creativity, many kinds of learning, most kinds of mental illness, and so on. Explanations in terms of belief, desire, sensation, and such tend not to work in these contexts.
 - D. Many folk psychological explanations face fairly direct empirical challenges. Split-brain experiments, for instance, suggest that we are good at convincing ourselves that we are acting for cogent reasons even when it's quite clear that we're not doing so.
 - E. The entities and laws posited by folk psychology do not cohere well with those of neuroscience and other parts of scientific psychology.

- F. The categories of folk psychology bring all sorts of problems in their wake, such as how we can believe in falsehoods, see or want non-existent things, and so on. Beliefs and the things believed are weird entities that tend to keep philosophers in business.
- III. If these considerations are on the right track (admittedly a big “if”), folk psychology would seem to be candidate for an *eliminative reduction*.
- A. In our earlier discussion of reduction, we found some reasons for respecting the explanatory power of theories that don’t seem to reduce to more basic theories. But that’s only true if the theories do a good job in their own domain. If not, they should be replaced by a better theory, in the same way that witchcraft is replaced by a theory that posits the existence of sexism, religion run amok, and some peculiar rules of legal procedure.
- B. The failure of fit with neuroscience figures prominently in arguments for the replacement of folk psychology. Many neuroscientists think that a *connectionist* model of the mind fits the scientific data much better than does the folk psychological theory.
1. On such a model, the brain does not “think” in states or episodes structured like sentences. Variations in patterns of stimulation across large numbers of neurons produce representations, just as variations in brightness levels produce an image on a television screen.
 2. Most information processing is, thus, subconceptual.
 3. Similarly, learning is less a matter of accumulating data stored as sentences than it is a matter of arranging stimulation patterns in the brain.
 4. The picture that emerges is not one according to which folk psychology is shallow and neuroscience deeper. The worry is that folk psychology presents a seriously misleading picture of what is really going on in our brains.
- IV. An eliminative reduction of folk psychology seems to have the consequence that none of us has beliefs, desires, and such and, hence, that none of us is a person, because we think of personality through the concepts of folk psychology.
- A. Is this really possible?
1. If folk psychology is anything resembling a scientific theory, we should be open to the possibility that it is largely false.
 2. It can seem problematic to deny the truth of folk psychology, because if it’s false, you can’t *believe* that it’s false, because on that hypothesis, there are no beliefs. But this problem can be surmounted rather readily.
- B. The best way out of the problem is probably to defend folk psychology as a decent semi-scientific theory.
1. The eliminativists sometimes stick folk psychology with problems that it needn’t face. It is not clear that it is folk psychology’s job to explain such phenomena as mental illness. It generally is used as a theory of normal, intelligent behavior.
 2. And some of folk psychology’s limitations are shared by its competitors. Folk psychology doesn’t have a good handle on creativity, but neither does neuroscience so far as I’ve been able to tell.
 3. Further, the concepts and ontology of folk psychology do seem to get used in serious and successful scientific psychology. Examples include rational choice theory, memory, and some parts of learning theory. In addition, folk psychology arguably plays a crucial role in some of the explanatory and predictive successes of history, economics, anthropology, and sociology.
 4. Perhaps most importantly, the worries about a failure of fit between folk psychology and neuroscience might be premature. We have a lot to learn in the whole domain of psychology.
- C. Folk psychology could be defended from elimination by denying that it is a competitor with scientific psychology. One could treat folk psychology in a strongly instrumentalist manner, for instance. We talk as if a thermostat has beliefs about the temperature in the house, but we aren’t being metaphysically serious about it. Might we be talking about each other that way?
- V. Surprisingly but importantly, issues of what theories are and what they are for, how they are confirmed, whether they can explain, how they fit with other theories, and so on are implicitly involved in our very sense of ourselves.

Essential Reading:

Churchland, "Eliminative Materialism and the Propositional Attitudes," in Boyd, Gasper, and Trout, *The Philosophy of Science*, pp. 615–630.

Supplementary Reading:

Greenwood, *The Future of Folk Psychology: Intentionality and Cognitive Science*.

Questions to Consider:

1. Sometimes quite a lot is at stake in a single contested term. Do you think that folk psychology should be understood as a psychological theory? Why or why not?
2. Imagine for a moment that you could become convinced that folk psychology is an irredeemably bad scientific theory. Would you give it up? What would that involve and what would be lost? To what extent is science answerable to our commonsense understanding of things and to what extent is our commonsense understanding of things answerable to science?

Lecture Thirty-Six

Philosophy and Science

Scope: In this lecture, I will attempt to forge some new connections among our by-now-old ideas. Our overarching themes involve tensions between tempting ideas: the distinctiveness of science versus the continuity between philosophy and science, the competing modesties of empiricism and realism, respect for accurate descriptions of scientific practice versus the legitimacy of attempts to improve the practice, and the importance of developing a picture of the intellectual virtues and values of science and scientists that is neither cynical nor smug. The lecture (and the course) aspires to leave you puzzled in articulate and productive ways.

Outline

- I. We began the course by wondering what is special about science. The idea that there is something distinctive about the sciences is attractive, but it sits awkwardly with attractive aspects of holism and naturalized epistemology.
 - A. It is not clear to what extent we want to distinguish scientific from everyday theorizing. We do not take ourselves to be doing science in our everyday lives, but we properly aspire to embody scientific virtues in many of our everyday undertakings.
 - B. Just as science had better be different from our everyday practices but not too different, science had better be different but not too different from philosophy. We've seen a number of reasons for thinking that science and philosophy should be thought of as continuous, but we mustn't lose sight of the distinctive manner in which science manages to put questions to nature.
 - C. The search for a demarcation criterion, however, does not look promising. Science probably cannot be done without some kind of metaphysical picture or conception lurking in the background.
 - D. The inescapability of metaphysics emerges most clearly in notions of categories, kinds, properties, and so on. What counts as two things being similar? Which terms can figure in laws? It's often hard to see what would count as an adequate defense of our category schemes.
 - E. How can these tensions between the distinctiveness and continuity of science be resolved or, at least, softened?
 1. We should recognize that our metaphysical views are not very readily tested by the world; thus, we should be modest, flexible, and self-conscious about them.
 2. We should think of science as differing from other pursuits in a number of medium-sized ways, rather than in one big way or in no way at all. Science involves a distinctive combination of observation, education, social structure, and other elements.
 - F. Much of the best philosophy of science these days reflects the continuity between philosophy and science because it is both informed and driven by empirical concerns. Quine's holism and naturalism help us to see that we are sometimes working on the same parts of the web from different angles. But this holism can make the task of distributing criticism across the web of belief challenging.
- II. Empiricism, both about meaning and about evidence, is an attractive idea, but it is difficult to keep it in check, and it sits poorly with scientific realism, which is an attractive idea.
 - A. Empiricism about meaning is particularly unfashionable with philosophers these days—and for good reason. It hamstring's our ability to talk about almost anything. But the lesson of special relativity still looms. The greater we extend our semantic search, the more risk we run of exceeding our epistemic grasp. Empiricism helps us avoid muddling our meanings.
 - B. Scientific realism seems compelling, but a realist needs to stay in touch with his or her inner empiricist. Inference to the best explanation is fragile under favorable conditions, and if we are going to be realists, we should squarely face the limitations of our evidential situation.
 - C. Realists need to remind themselves of the epistemic risks they run and should try to be as clear as possible about the intellectual benefits of the risks. For their part, empiricists need to remind themselves of the

intellectual (for example, explanatory) resources of which their empiricism deprives them and should strive to be clear about the benefits of their relative asceticism.

1. If you stick closely to what is given in experience (and do not assume that too much is given), you will avoid certain mistakes. Popper's skepticism about induction falls into this camp, as does resistance to Bayesian subjective probabilities. But increased security comes at the cost of diminished resources and an increased vulnerability to skepticism.
2. On the other hand, such things as explanatory ambitions, models, and analogies and a willingness to take subjective probability assignments seriously allow one to set out to maximize the range and depth of one's beliefs. But here the risks are muddleheadedness and mistakes.

III. Kuhnian fidelity to actual science is an attractive idea, but it sits awkwardly with the "is/ought" distinction.

- A. The smart money is on scientific practice over philosophical advice about scientific method. But what scientists say they're doing doesn't always reflect what they're actually doing. Scientists tend to commit philosophy when they explain what they do. And we've seen a number of respects in which sympathetic observers might think that scientific practice might be improved.
- B. Such terms as *objectivity* can be dangerous because they tend to lead people to exaggerate the virtues and/or the vices of science and scientists.
 1. People sympathetic to such views as social constructivism hear talk of objectivity and picture scientists claiming to view nature from nowhere, to step out of their skins, to carve nature at its joints, and so on. They rightly regard most of this as pretty naïve, but that stems from the too-demanding notion of objectivity being deployed.
 2. People sympathetic to realism and/or empiricism hear scoffing about objectivity and think that science is being reduced to mere rhetoric or worse. They then tend to dismiss legitimate questions about, for example, the role of values in science. This leads to soft-headed thinking about such things as the problem of demarcation.
- C. Science deserves a distinctive kind of respect. No amount of examining it, warts and all, undermines its achievements. But we laypeople should not accord it automatic deference.

IV. Philosophy, especially philosophy of science, is hard. It compensates us only with clarity, with the ability to see that the really deep problems resist solutions. But clarity is not such cold comfort after all. As Bertrand Russell argued, it can be freeing. When things go well, philosophy can help us to see things and to say things that we wouldn't have been able to see or to say otherwise.

Essential Reading:

Godfrey-Smith, *Theory and Reality: An Introduction to the Philosophy of Science*, chapter 15.

Supplementary Reading:

Rosenberg, *Philosophy of Science: A Contemporary Introduction*, chapter 7.

Questions to Consider:

1. Which has changed more as a result of this course, your conception of science or your conception of philosophy?
2. Which of the tensions sketched in the last lecture do you think it would be desirable to resolve and which ones seem essential to the success of the scientific enterprise (and, hence, need to be left unresolved)? Can we, when doing serious intellectual work, simultaneously treat something as a puzzle or problem yet think it is better left unsolved?

Biographical Notes

Ayer, Alfred Jules (1910–1989). Ayer made a splash as a young man when he published *Language, Truth and Logic* in 1936. That book provides the classic statement in English of logical positivism. Ayer remains best known for this brash, youthful book, but he went on to do important work in several areas of philosophy. He also made a record with Lauren Bacall!

Berkeley, George (1685–1753). Berkeley was born near Kilkenny, Ireland. He became an ordained Anglican minister in 1710 and was appointed bishop of Cloyne in 1734. His first important philosophical work concerned the theory of vision, and he later incorporated results into the God-centered, immaterialist conception of the world defended in his most important book, *A Treatise Concerning the Principles of Human Knowledge*, published in 1710. Alexander Pope said that Berkeley possessed “ev’ry virtue under heav’n.”

Bridgman, Percy (1882–1961). Bridgman was born in Cambridge, Massachusetts; attended Harvard; and spent his academic career there. He received the Nobel Prize for Physics in 1946 for his work on the properties of materials subjected to high pressures and temperatures. Bridgman’s experimental work has proved important in geology and for such processes as the manufacture of diamonds. His most important work in the philosophy of physics is *The Logic of Modern Physics*, which was published in 1927.

Carnap, Rudolf (1891–1970). Born and educated in Germany, Carnap joined the Vienna Circle in the 1920s. In 1928, he published *The Logical Structure of the World*, an ambitious attempt to reduce talk of objects and such to experiential terms. He made a number of major contributions to philosophy in the logical positivist tradition, including crucial work on inductive logic and the structure of scientific theories. Carnap came to the United States in 1935 and spent most of his academic career at the University of Chicago and at UCLA.

Einstein, Albert (1879–1955). Nothing in Einstein’s early career would have led anyone to expect his *annus mirabilis*, which took place in 1905. During that year, he published the essentials of special relativity and did groundbreaking work on Brownian motion and the photoelectric effect. He completed his general theory of relativity in 1915, and when that theory received impressive confirmation from the eclipse experiment of 1918, Einstein became an international celebrity. Einstein spent much of the rest of his scientific career pursuing a grand unified theory, and he also made important contributions to the philosophy of science, all the while crusading for peace.

Feyerabend, Paul (1924–1994). Feyerabend was born in Vienna just as the Vienna Circle was coming together. He was shot in the spine in 1945 while serving in the German army. After studying singing, history, sociology, and physics, he wrote a philosophy thesis and went to England to study with Karl Popper. Feyerabend’s critiques of the dominant empiricist accounts of observation and reduction culminated in his rejection of the whole idea of scientific method, as most influentially expressed in his 1975 book, *Against Method*. Late in his career, Feyerabend spent much of his time articulating and defending philosophical relativism. Most of his academic career was spent at the University of California at Berkeley.

Goodman, Nelson (1906–1998). Goodman was born in Massachusetts and educated at Harvard. Before beginning his teaching career, he was the director of a Boston art gallery. He taught at the University of Pennsylvania and at Brandeis before joining the Harvard faculty in 1968. Though he is best known for his “new riddle of induction” (also known as the “grue problem”), Goodman made important contributions to aesthetics, philosophy of language, and epistemology, as well as philosophy of science. His 1978 book *Ways of Worldmaking* probably provides the best introduction to his distinctive approach to philosophical questions.

Hempel, Carl Gustav (1905–1997). Born in Orianenburg, Germany, Hempel studied logic, mathematics, physics, and philosophy at several German universities. He was a member of the Berlin Circle of logical positivists before moving to Vienna to work with members of the Vienna Circle. After coming to the United States in 1939, Hempel taught at Queens College, Yale University, Princeton University, and the University of Pittsburgh. Hempel’s covering-law approach to explanation dominated the field for decades, and he made important contributions to the theory of confirmation, as well. His introductory text, *Philosophy of Natural Science* (1966), is regarded as a classic and offers a clear and readable approach to the field.

Hume, David (1711–1776). Often considered the greatest of the empiricist philosophers, Hume was born in Edinburgh. His *A Treatise of Human Nature*, written while Hume was in his 20s, is now regarded as one of the great

works of modern philosophy but was largely ignored at the time. Doubts and whispers about Hume's religious views prevented him from ever attaining an academic position in philosophy. Hume's six-volume *History of England* did provide him with a good measure of literary success, however. His posthumously published *Dialogues Concerning Natural Religion* is generally considered a masterpiece. Hume counted Adam Smith among his good friends, and he befriended Jean-Jacques Rousseau, though they later had a very public falling out.

Kuhn, Thomas S. (1922–1996). Born in Cincinnati, Ohio, Kuhn did his undergraduate work at Harvard and received his Ph.D. in physics from the same institution in 1949. By that point, however, he had developed serious interests in the history and philosophy of science. Kuhn began his teaching career at Harvard before moving to Berkeley, Princeton, and, finally, to M.I.T. *The Structure of Scientific Revolutions* (first published in 1962) made such terms as *paradigm* and *incommensurability* part of everyday academic discourse. The book remains enormously influential. Though *Structure* issued serious challenges to the picture of science as unproblematically progressive, cumulative, and objective, Kuhn himself saw science as an unrivaled epistemic success story.

Lakatos, Imre (1922–1974). Lakatos was born and raised a Jew, though he later converted to Calvinism. His mother and grandmother died in Auschwitz. He worked in a Marxist resistance group during the Nazi occupation of his native Hungary. The communist government after the war placed him in an important position in the Ministry of Education, but he was arrested for “revisionism” in 1950. He spent almost four years in prison, including a year in solitary confinement. Lakatos took a leadership role in the Hungarian uprising of 1956 and left his native country after the Soviets suppressed the rebellion. His Ph.D. dissertation (written at Cambridge University) eventually became *Proofs and Refutations*, a remarkable work in the philosophy of mathematics. After receiving his doctorate, Lakatos joined the Popper-dominated philosophy department at the London School of Economics, where he remained until his premature death. He spent the bulk of his career developing and defending his *methodology of scientific research programs*, in which he tried to combine a Kuhnian historicism with an objective methodological standard.

Locke, John (1632–1704). One of the most influential philosophers of the modern period, Locke was educated at Westminster School in London and at Christ Church, Oxford. While at Oxford, he studied and eventually taught logic, rhetoric, and moral philosophy. He also became interested in the relatively new experimental and observational approach to medicine. In 1667, Locke moved to London as the personal physician, secretary, researcher, and friend of Lord Ashley. Lord Ashley eventually became the First Earl of Shaftesbury and Lord Chancellor, and through him, Locke became deeply involved in the turbulent politics of the period. Locke's most important work, *An Essay Concerning Human Understanding* (1690), stems in part from political motives, as Locke hoped to determine which questions could be addressed by human reason so that fruitless debates could be avoided. Other important works, including *Two Treatises of Government* (1690) and *Letter Concerning Toleration* (1690), more directly reflect Locke's concern with public life.

Mill, John Stuart (1806–1873). Mill's father, James, was himself an important philosopher, and he gave his son a remarkably intense education (Mill began reading Greek at the age of 3). His rigorous childhood left the younger Mill intellectually precocious but emotionally stunted, and he suffered a debilitating “mental crisis” in his early 20s. An exposure to the arts helped him overcome his depression. Mill is probably best known as a moral and political philosopher; *Utilitarianism* (1863), *On Liberty* (1859), and *The Subjection of Women* (1869) are classics in those fields. Mill spent much of his life working for the East India Company and was a member of Parliament from 1865 to 1868. His thoroughgoing empiricism in epistemology and metaphysics emerges in his *System of Logic* (1843) and *Examination of Sir William Hamilton's Philosophy* (1865). Mill's *Autobiography* (1873) is also a classic.

Popper, Karl (1902–1994). Popper grew up in Vienna and took his Ph.D. from the University of Vienna in 1928. He shared many scientific and philosophical interests with the members of the Vienna Circle but disagreed with them enough that he was not invited to become a member. Popper's *Logic of Scientific Discovery* (1934) presented his own falsificationist, anti-inductive conception of scientific inquiry, along with his criticisms of logical positivism. The book remains influential to this day. The rise of Nazism forced Popper to flee to the University of Canterbury in New Zealand, where he turned his attention to social and political philosophy. *The Poverty of Historicism* (1944) and *The Open Society and Its Enemies* (1945) are products of that period. Popper moved to the London School of Economics in 1949 and was knighted in 1964.

Quine, Willard van Orman (1908–2000). Quine was born in Ohio and attended Oberlin College. He did his graduate work at Harvard and spent his academic career there. Soon after receiving his Ph.D., Quine traveled to Vienna, where he worked with the leading positivists, and to Prague, where Carnap was then living. Carnap made

an enormous impression on Quine, and much of Quine's work can usefully be seen as responding to the problems faced by Carnap's versions of positivism and post-positivism. Quine's many influential papers in the philosophy of language, logic, philosophy of mind, and science are scattered through a number of collections.

van Fraassen, Bastian (1941–). Born in the Netherlands and educated at the University of Alberta and University of Pittsburgh, van Fraassen has spent much of his career working out what empiricism can amount to after the demise of positivism. He has taught at Yale University, the University of Toronto, and the University of Southern California, and he has been on the faculty at Princeton since 1982. His most influential work in the philosophy of science is *The Scientific Image* (1980), and he has also done important work in philosophical logic.

Bibliography

Essential Reading: General Anthologies

Balashov, Yuri, and Alex Rosenberg, eds. *Philosophy of Science: Contemporary Readings*. New York: Routledge, 2002. This anthology works particularly well with Rosenberg's textbook (see below). Like the Rosenberg text, it is organized rather differently than our course, but it does an especially nice job of finding fairly accessible writings that nevertheless touch on the crucial issues.

Boyd, Richard, Philip Gasper, and J. D. Trout, eds. *The Philosophy of Science*. Cambridge, MA: MIT Press, 1991. This anthology has the virtue of bringing in a good bit of material from the philosophy of particular sciences. As a result, its coverage of general philosophy of science is a bit spottier than that of the other anthologies listed here, but it still does quite a nice job.

Curd, Martin, and J. A. Cover, eds. *Philosophy of Science: The Central Issues*. New York: W.W. Norton & Co., 1998. This is the anthology I use in my courses. I find it shockingly light on coverage of logical positivism, but the readings are, for the most part, otherwise well chosen, and this anthology includes extensive and user-friendly commentary from the editors. Given that even introductory anthologies in the philosophy of science consist of pretty difficult material, the commentary is especially useful.

Hitchcock, Christopher, ed. *Contemporary Debates in Philosophy of Science*. Malden, MA: Blackwell Publishing, 2004. Unlike the other books listed in this section, this is not really a comprehensive anthology, though it does cover a reasonable range of issues. It is organized as a series of debates, and it can be illuminating to see professional philosophers engaged in direct disagreement. Most of the contributions are both accessible and lively.

Essential Reading: General Textbooks

Godfrey-Smith, Peter. *Theory and Reality: An Introduction to the Philosophy of Science*. Chicago: University of Chicago Press, 2003. This is my favorite introductory text. Godfrey-Smith's sense of how to organize this material accords with mine, and his writing is clear and accessible. Because our course covers more material and does so in more depth than this book does, the text might appear a bit slight after working through our course, but it would make a nice accompaniment to it.

Hung, Edwin. *The Nature of Science: Problems and Perspectives*. Belmont, CA: Wadsworth Publishing, 1997. A thorough and accessible textbook that gives special attention to the study of patterns of reasoning that operate in science. It covers a lot of ground twice (once topically and once historically), which can be illuminating if not terribly efficient.

Rosenberg, Alex. *Philosophy of Science: A Contemporary Introduction*. London: Routledge, 2000. A nice introduction that simply approaches the material in a different order than I do. For this reason, this might prove a particularly useful book to some who have heard this course; it provides a different way of seeing how this material hangs together. Rosenberg focuses somewhat more narrowly than do the other authors considered here.

Essential Reading: Particular Topics

Ayer, Alfred Jules. *Language, Truth and Logic*, 2nd ed. New York: Dover Publications, 1952. The great positivist manifesto, at least in the English language. This work is enlivened by confidence and anti-metaphysical fervor. It presents positivism as a philosophical program, touching on philosophy of language, metaphysics, and ethics, as well as philosophy of science.

Berkeley, George. *Three Dialogues between Hylas and Philonous*. New York: Oxford University Press, 1998. There are many editions of this wonderful book, and just about any of them will do. See whether you can fare better than poor Hylas does as Philonous (speaking for Berkeley) attacks the idea that matter is a useful, a necessary, or even an intelligible concept.

Bird, Alexander. *Thomas Kuhn*. Princeton: Princeton University Press, 2000. Kuhn's ideas are of enormous philosophical importance, but he wasn't a philosopher; thus, it's quite useful to have an accessible, book-length treatment by a philosopher. For our purposes, one especially nice feature of this book is its emphasis on the positions Kuhn shares with his positivist predecessors.

Feyerabend, Paul. *Against Method*, 3rd ed. London: Verso, 1993. A provocative polemic against the idea of science as a rule-governed enterprise. Feyerabend celebrates the creative and anti-authoritarian side of science. He pushes

pluralism a bit further than most philosophers think it can go, but the book is refreshing and contains important lessons.

Greene, Brian. *The Elegant Universe: Superstrings, Hidden Dimensions, and the Quest for the Ultimate Theory*, 2nd ed. New York: Vintage Books, 2003. An astonishingly accessible introduction to the strangeness that has dominated physics for the past century. Superstring theory is the main topic of the book, and I barely touch on that in this course, but along the way, Greene offers vivid and undemanding introductions to relativity and quantum mechanics, both of which make appearances in our course.

Hacking, Ian. *The Emergence of Probability*. Cambridge: Cambridge University Press, 1975. A fascinating story about how and why the central mathematical ideas of probability and statistics emerged in Europe in the 17th century. The book is mathematically undemanding and sheds light on economics, theology, and gambling, as well as history and philosophy.

Hume, David. *A Treatise of Human Nature*. New York: Oxford University Press, 2000. Hume's most important philosophical work, originally published while he was in his 20s. From the standpoint of our course, it is Hume's unwavering commitment to empiricism that matters most, but his unrepentant naturalism is of equal philosophical importance. This edition has some valuable editorial material, but other editions are more than adequate to the purposes of our course.

Kuhn, Thomas S. *The Structure of Scientific Revolutions*, 3rd ed., enlarged. Chicago: University of Chicago Press, 1996. The influence of this book can be overestimated, but it isn't easy to do so. *Structure* is often maddeningly unclear if considered as philosophy (rather an unfair standard to apply to a history book), but it decisively established the philosophical importance of the history of science and revolutionized the study of science from several disciplinary perspectives.

Sterelny, Kim, and Paul E. Griffiths. *Sex and Death: An Introduction to Philosophy of Biology*. Chicago: University of Chicago Press, 1999. How good a title is that? A fun and thorough introduction to the philosophy of biology, very generously laced with examples. I draw on only one chapter of it in this course, but I recommend the book as a whole. Readers should be warned, however, that *Sex and Death* is probably less even-handed and uncontroversial than some people expect introductory texts to be.

Supplementary Reading

Brody, Baruch, and Richard E. Grandy, eds. *Readings in the Philosophy of Science*, 2nd ed. Englewood Cliffs, NJ: Prentice Hall, 1989. This very useful anthology (full disclosure: I worked with Brody and Grandy when I was an undergraduate) has recently gone out of print, but as of this writing is easily and cheaply available through such outlets as amazon.com. It remains a convenient way to acquire some classic articles.

Carroll, John W., ed. *Readings on Laws of Nature*. Pittsburgh: University of Pittsburgh Press, 2004. Another case requiring full disclosure: Carroll is a colleague of mine at N.C. State. But I'm not alone in my opinion that he knows as much about laws of nature as anybody alive. As is to be expected in an anthology on a relatively specialized topic in the philosophy of science, the water gets a bit deep here, but this is as wide-ranging and accessible a collection of essays on laws of nature as one could ask for.

Clark, Peter, and Katherine Hawley, eds. *Philosophy of Science Today*. Oxford: Clarendon Press, 2000. An expanded version of the 50th-anniversary issue of the *British Journal for the Philosophy of Science*, this collection is aimed at professional philosophers, but the essays are surveys and some of them are fairly accessible. The collection as a whole provides a nice summary of the current state of the field.

Greenwood, John D., ed. *The Future of Folk Psychology: Intentionality and Cognitive Science*. Cambridge: Cambridge University Press, 1991. This volume stands at the intersection of philosophy of science and philosophy of mind. Given that we haven't learned much philosophy of mind in this course, some of the essays collected here are forbidding. But the philosophy of mind literature hasn't gotten any easier since this collection came out, and this remains a convenient place from which to survey debates about the relationships between folk psychology and scientific psychology.

Hacking, Ian. *The Taming of Chance*. Cambridge: Cambridge University Press, 1990. A follow-up to *The Emergence of Probability*, this book continues Hacking's engaging and interdisciplinary story through the 19th century. This volume tells the story of how chance came to make the world seem orderly.

Kitcher, Philip. *Science, Truth, and Democracy*. Oxford: Oxford University Press, 2001. This seems to me an important book. It begins by tackling issues about truth and reality, proceeds through some difficult issues about

objectivity and interests, and culminates in a sure-to-be-influential proposal concerning the proper function of science in a democracy.

Kornblith, Hillary, ed. *Naturalizing Epistemology*, 2nd ed. Cambridge, MA: MIT Press, 1994. An influential and reasonably accessible collection of papers marking the “naturalistic turn” in epistemology and philosophy of science. Some of the essays are examples of naturalized epistemology, while others confront the philosophical issues about circularity and normativity that are highlighted in our course.

Ladyman, James. *Understanding Philosophy of Science*. London: Routledge, 2002. This is another nice introductory textbook, and it is especially generous in its explanations of the philosophical problems surrounding inductive inference. This book offers a somewhat narrow but impressively clear and helpful introduction to our field.

Larvor, Brendan. *Lakatos: An Introduction*. London: Routledge, 1998. I confess to feeling a bit guilty about including this work, rather than something by Lakatos himself, on the reading list. But Larvor is easier to read than is his subject, and the reader gets the added benefit of having Lakatos’s marvelous work in the philosophy of mathematics presented, along with his contributions to the philosophy of science.

Nickles, Thomas, ed. *Thomas Kuhn*. Cambridge: Cambridge University Press, 2003. Kuhn’s work has elicited an astonishing range of reactions, and it bears on a great many disciplines. This collection of essays by 10 different thinkers on 10 different topics is a good way to take the measure of Kuhn’s legacy.

Pennock, Robert, ed. *Intelligent Design Creationism and Its Critics: Philosophical, Theological, and Scientific Perspectives*. Cambridge, MA: MIT Press, 2001. A collection edited by an unabashed critic of intelligent-design creationism, but one that lets advocates of intelligent design speak for themselves. Though Pennock is a philosopher of science, the collection ranges across history, politics, law, biology, theology, and education.

Shapin, Steven, and Simon Schaffer. *Leviathan and the Air-Pump: Hobbes, Boyle, and the Experimental Life*. Princeton: Princeton University Press, 1985. Like many philosophers of science, I have some serious qualms about the notions of truth, reality, and construction at work in much recent sociology of science. But this fascinating story about the history and politics behind the rise of the experimental method more than compensates the reader for any lingering philosophical irritations.

Sklar, Lawrence. *Philosophy of Physics*. Boulder, CO: Westview Press, 1992. This one is not exactly light reading, but it offers a clear and effective presentation of the deep philosophical puzzles that arise about space and time, quantum mechanics, and the role of probability in physics. Sklar does a lovely job of using the philosophy and the physics to illuminate each other.

Soames, Scott. *Philosophical Analysis in the Twentieth Century*, 2 vols. Princeton: Princeton University Press, 2005. This enormously useful survey centers more on metaphysics and the philosophy of language than on the philosophy of science, but this is a terrific way to learn about the developments in philosophy of language that made logical positivism possible. The later parts of the story illuminate Quinean holism and the theory of reference that helped make way for the resurgence of scientific realism.

Sober, Elliott. *Philosophy of Biology*, 2nd ed. Boulder, CO: Westview Press, 2000. A straightforward and sophisticated introductory text by one of the leading contemporary philosophers of biology.

Woolhouse, R. S. *The Empiricists*. New York: Oxford University Press, 1988. A brisk, reliable survey, not just of Locke, Berkeley, and Hume, but of the empiricist ideas at work in such predecessors as Bacon and Hobbes. This book is very friendly to beginners in philosophy.

Reference Works

Machamer, Peter, and Michael Silberstein, eds. *The Blackwell Guide to the Philosophy of Science*. Malden, MA: Blackwell Publishers, 2002. This is a very handy book. It consists of article-length essays, most of them about as accessible as their subject matters permit, that survey a topic or problem and offer suggestions about future developments. The volume touches on classic problems, such as explanation, but also on less common topics, such as metaphor and analogy in science.

Newton-Smith, W. H., ed. *A Companion to the Philosophy of Science*. Malden, MA: Blackwell Publishers, 2000. The *Companion* works differently than the *Guide* put out by the same publisher. This book is even more useful than its cousin. It offers short essays on dozens of topics and figures. Each entry provides the bare essentials concerning its subject, and most of the entries are quite accessible indeed.

Internet Resources

Philosophy of Science Resources. <http://pegasus.cc.ucf.edu/~janzb/science>. A terrific source of information, this page contains a great many links and construes the philosophy of science broadly.

Philosophy of Science Undergraduate Research Module at N.C. State.

<http://www.ncsu.edu/project/ungradreshhmi/evaluationModule/login.php>. This mini-course was prepared by three of my colleagues at N.C. State (before I joined the department). The module covers demarcation, confirmation, explanation, and reduction. The writing is clear and helpful. The site requires users to obtain a password and to log in, but the site is free and requires no membership. The log-in enables granting agencies to track how many people use the module.

Science Timeline. <http://www.sciencetimeline.net>. An impressively detailed timeline that includes a surprising number of philosophical references and makes it easy to obtain more detailed information.

The Stanford Encyclopedia of Philosophy. <http://plato.stanford.edu>. This is a marvelous and growing peer-reviewed reference site for all of philosophy. Accordingly, not everything at the site bears on our course, and not every entry is accessible to non-specialists. But the site is easy to navigate and many of the entries are quite accessible.