Lee McIntyre Author of Post-Truth

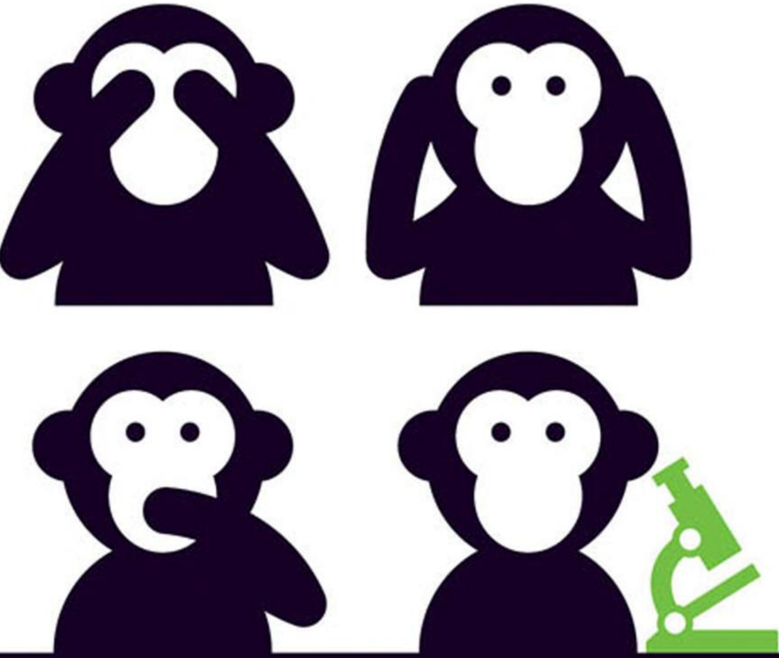

The Scientific Attitude

Defending Science from Denial, Fraud, and Pseudoscience

The Scientific Attitude

Defending Science from Denial, Fraud, and Pseudoscience

Lee McIntyre

The MIT Press
Cambridge, Massachusetts
London, England

This book was set in Stone Serif by Westchester Publishing Services. Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Names: McIntyre, Lee C., author.
Title: The scientific attitude : defending science from denial, fraud, and pseudoscience / Lee McIntyre.
Description: Cambridge, MA : The MIT Press, [2019] | Includes bibliographical references and index.
Identifiers: LCCN 2018037628 | ISBN 9780262039833 (hardcover : alk. paper)
Subjects: LCSH: Science—Social aspects. | Science—Methodology. | Pseudoscience.
Classification: LCC Q175.5 .M3955 2019 | DDC 306.4/5—dc23
LC record available at https://lccn.loc.gov/2018037628

10 9 8 7 6 5 4 3 2 1

Contents

Preface

This book has been a labor of love from the beginning and—as with any labor—it has taken a while. I remember the exact moment when I decided to become a philosopher of science, as I was reading Karl Popper's enchanting essay "Science: Conjectures and Refutations," in the fall of 1981, in one of the upper carrels in Olin Library at Wesleyan University. The issues seemed earth-shattering and the romance was obvious: here was a person who had found a way to defend one of the ideas I believed in most—that science was special. Popper made it his life's work to defend the epistemic authority of science and explain why it was superior to its imposters. How could I not want to be involved in that?

Though the issues gripped me, I never fully agreed with Popper's conclusions. I knew I'd get back to it someday but, as the professional reward system in academia seemed to favor taking somewhat smaller bites of the apple, I contented myself with spending the first decade of my career writing about the importance of laws and prediction, how to improve the methodology of the social sciences, and why we needed a philosophy of chemistry. Since then I have taken great enjoyment in expanding my reach to write philosophy for a general audience on topics such as science denial, the importance of reason, and why—especially in this day and age—even the staunchest philosophical skeptics need to defend the idea that truth matters.

But this is the book that I have always wanted to write. By taking up a topic as important as what is distinctive about science, I hope that it will be of interest to philosophers, scientists, and the general public alike.

For inspiring me to go into philosophy I would like to thank my teachers: Rich Adelstein, Howard Bernstein, and Brian Fay. Although I overlapped with him only toward the end of my college career, Joe Rouse was also an inspiration to me. In graduate school at the University of Michigan, I had the good fortune to learn the philosophy of science from Jaegwon Kim, Peter Railton, and Larry Sklar. I was not always happy in graduate school (who is?), but I look back on my education there as the foundation for all my further work.

Since then, I am grateful to have worked with some of the best in the business: Dan Little, Alex Rosenberg, Merrilee Salmon, and Eric Scerri, all of whom have taught me much through their excellent scholarship and warm colleagueship. My debt to Bob Cohen and Mike Martin—both of whom passed away in recent years—is enormous, for they gave me a home in the philosophy of science and helped me along at every step of the way. I am glad to say that the same has been true of the new director of the Center for Philosophy and History of Science at Boston University, Alisa Bokulich, as well.

For guidance and advice on some of the specific ideas contained in this book, I would like to thank Jeff Dean, Bob Lane, Helen Longino, Tony Lynch, Hugh Mellor, Rose-Mary Sargent, Jeremy Shearmur, and Brad Wray. I was lucky to have been a participant in Massimo Pigliucci and Maarten Boudry's workshop on "scientism" at CUNY in the spring of 2014, where I heard some extremely stimulating papers by Noretta Koertge, Deborah Mayo, and Tom Nickles, that inspired me to think about writing this book. Rik Peels and Jeff Kichen made pinpoint suggestions about discrete problems that were enormously helpful as well.

My good friends Andy Norman and Jon Haber have done me the honor of reading the complete manuscript of this book in draft and making many helpful suggestions. My friend Laurie Prendergast has once again done yeoman service by helping me with the proofreading and index. I would also like to acknowledge the work of five anonymous referees, whom I obviously cannot thank by name, each of whom made enormous contributions to the content of this book. It goes without saying that any remaining mistakes are mine alone.

My father unfortunately did not live to see the publication of this book, but to him and to my mother I send all my love and gratitude for always believing in me and for their support and guidance over the years. My wife Josephine, and children Louisa and James, each read this book in detail and lived with its ups and downs through many iterations. No man was ever luckier to be married to such a wonderful woman, who wants nothing more than for me to be happy in my life and in my work. I am fortunate also to have not one but *two* children who majored in philosophy and claim it as their birthright to identify any flaws in the old man's reasoning, which they have done with frightening efficiency. Indeed, my children's contributions to this book have been so great that I would like to dedicate it to them.

The team at MIT Press is without parallel. As they proved in my last book with them—and every day since—no author is ever successful alone. From copyediting to design, and marketing to editorial, it is my privilege to work with all of them. Here I would like to give special thanks to my tireless and creative publicity team and to my editor Phil Laughlin, who is analytical, succinct, practical, and hilarious all at the same time. Once again they have made it a pleasure to work with MIT Press, in what is now my fourth book with them.

My final debt is an old one, but I see it every day as I look at the framed handwritten letter that I received from Karl Popper in March 1984, in response to a letter I wrote to him as an undergraduate. Popper was brilliant, lucid, defensive, and enlightening. Although I disagree with many of his ideas in the philosophy of science, I could not have developed my own without having had his to react against and—in one of the most delightful discoveries of my career—found that in some ways he had already anticipated the scientific attitude. I never met Karl Popper, but my earliest vision of him stays with me: a young man just starting out in the winter of 1919, realizing the logic of falsification in a lightning flash, then working out its details over the course of his career. I was proud to learn that this book would be published precisely one hundred years after Popper's discovery. It is small tribute to a man who inspired my own and so many other careers in the philosophy of science.

whether it is even a worthy aim. The idea that we can transplant science into other fields by understanding what is most distinctive about it has gotten something of a bad reputation over the years. This notoriety has come from those who have claimed that there is a "scientific method"—or some other firm criterion of demarcation between science and nonscience—such that if we could just apply the standard rigorously enough, good science would bloom as a result. Such claims are made worse by those who embrace the spirit of proselytizing and engage in what has been called "scientism," whereby they now have a hammer and every other field in the universe of inquiry looks like a nail. But there is a problem: nearly everyone in the philosophy of science these days admits that there is no such thing as scientific method, that trying to come up with a criterion of demarcation is old-fashioned, and that scientism is dangerous.[7] Along the way, most have also largely given up on the idea that prescription lies at the heart of the philosophy of science.

Karl Popper's 1934 model of science *Logik der Forschung*, translated into English in 1959 by Popper himself as *The Logic of Scientific Discovery*, focuses heavily on the idea that there is a reliable method for demarcating science from nonscience, but that there is no such thing as a "scientific method." Popper champions the idea that science uses "falsifiable" theories—ones that are capable at least in principle of being proven wrong by some evidence—as the dividing line. Although this model has several logical and methodological virtues, it has also proven problematic for many philosophers of science, who complain that it is too idealized and focuses too heavily on the "greatest moments" of science, like the transition from the Newtonian to the Einsteinian model in physics, and that most science does not actually work like this.[8]

Another account was offered by Thomas Kuhn in 1962, in his famous book *The Structure of Scientific Revolutions*. Here the focus is on how some scientific theories replace others through paradigm shifts, where the scientific consensus changes radically as a result of problems that have built up with the old theory, and the field switches seemingly overnight to a new one. The problem here, however, is not only the familiar complaint that most science does not actually work like this (for instance, the transition from the Ptolemaic Earth–centered model to the Copernican Sun–centered model in astronomy)—which Kuhn freely admits when he talks about the ubiquity of "normal science"—but that even where it does, this is not a completely "rational" process. Although Kuhn insists on the key role of evidence in paradigm change, once we have opened the door to "subjective" or social factors in interpreting this evidence, there seems to be no "method" to follow.[9] This not only presents a problem for showing that scientific claims are justifiable, but also forestalls delineating a roadmap for other sciences.

Still further models of scientific change have been proposed by Imre Lakatos, Paul Feyerabend, Larry Laudan, and the "social constructivists," each of whom drain a little more water from the pool that allows us to say that science is "special" and that other fields of inquiry would do well to follow its example.[10] So what to do? Just pick one of the existing accounts? But this is not possible. For one thing, they are largely incompatible with one another; each describes a different piece of the "blind man's elephant," so we are still missing a comprehensive picture of what science is like. Another problem is that these models seem to succeed only to the extent that they leave something behind, namely the motivation that if we finally understood science we could provide a standard for other fields to become more scientific.

If all of the best accounts fail, perhaps there is a more general weakness in this whole approach? Although some may be loath to identify it as a weakness, it seems at least a drawback that the philosophy of science has spent so much of its time focusing on the "successes" of science and has not had very much to say about its failures. In fact, the lessons of failure to live up to the scientific standard are as revealing about what science is as the example of those fields that have achieved it. There is nothing wrong in principle with exploring what is distinctive about science by looking at its accomplishments, but this has led to some mischief.

First, while it would be comforting to imagine science as a long series of steps toward the truth—with its failures due only to the wrongheaded and ignorant—this view of science is belied by its history, which is littered with theories that were scientific but just turned out to be wrong. Both Popper and Kuhn have done much to show how science is strengthened by an uncompromising focus on explanatory "fit" between theory and evidence, but it is all too easy for others to look back and pretend that this was all inevitable and that the arc of science bends irrevocably toward a single (true) explanation of reality.

Second, the relentless focus on explaining science through its successes has meant that most of the "victories" that philosophers can turn to for their models have come from the natural sciences. Specifically, we have been forced to draw most of our conclusions about what makes science special from the history of physics and astronomy. But this is a bit like drawing targets around where the darts have landed. And does this mean that in attempting to be scientific, other fields should try to emulate physics? Thinking that the answer to this question is an unqualified yes has done a great disservice to other fields, some of which are solidly empirical but quite different in subject from the physical sciences. Remember that one important part of the mission of the philosophy of science is to understand what is distinctive about science *so that we can grow it elsewhere*. But where does this leave fields like the social sciences, which until quite recently have been underserved by most of the explanatory models in the philosophy of science?

Popper famously argued that the social sciences could not be sciences because of the "open systems" problem created by the effect of free will and consciousness on human decision making. In natural science, he claimed, we use falsifiable theories, but this path is not open to the social sciences.[11] Similarly, Kuhn, for all his fans in the social sciences (who felt that they may finally have a target they could hit), also tried to distance himself from the messy study of human behavior, by insisting that his model was applicable only to the natural sciences, and that he was not providing any advice to the social sciences. Add to this the problem of what to do about some of the other "special" (i.e., nonphysical) sciences—like biology or even chemistry—and we have a full-fledged crisis on our hands in defending a view of science that is separate from reduction to physics. What to do about the claim that there are epistemically autonomous concepts in chemistry (such as *transparency* or *smell*)—just as there are in sociology (such as *alienation* or *anomie*)—that cannot be explained at the physical level of description? If our model of successful science is physics, will even chemistry make the cut? From a certain perspective, *most* of those fields that are either scientific, or wish to become so, do not fit the models of philosophy of science that have been drawn from the history of physics and astronomy, and so could be considered "special sciences." Have we no advice, or justification, to offer them?

Finally, what to say about those fields that make a claim to being scientific, but just do not measure up (such as "intelligent design theory" or denialism about climate change)? Or of those instances where scientists have betrayed their creed and committed fraud (such as Andrew Wakefield's work purporting to show a link between vaccines and autism)? Can we learn anything from them? I maintain that if we are truly interested in what is special about science, there is much to learn from those who have forsaken it. What is the proponent of intelligent design theory *not* doing that genuine scientists *should* do (and in fact generally succeed in doing)? Why are climate change deniers unjustified in their high standards of "skepticism"? And why is it forbidden for scientists to rig their data, cherry pick their sample sets, and otherwise try to fit the data to their theory, if they want to succeed in scientific explanation?[12] It may seem obvious to those who defend and care about science that all of the above have committed a mortal sin against scientific principles, but shouldn't this help us in articulating the nature of those principles?

In this book, I propose to take a very different approach from my predecessors, by embracing not only the idea that there is something distinctive about science, but that the proper way to understand it is to eschew the exclusive focus on the successes of natural science. Here I plan to focus on those fields that have failed to be scientific, as well as those (like the social sciences) that might wish to become more so. It is one thing to discern what is distinctive about science by examining the transition from Newton to Einstein; it is another to get our boots muddy in the questions of scientific fraud, pseudoscience, denialism, and the social sciences.

Why bother? Because I think that to truly understand both the power and the fragility of science we must look not just at those fields that are already scientific, but also at those that are trying (and perhaps failing) to live up to the standard of science. We can learn a lot about what is special about science by looking at the special sciences. And we should be prepared to answer the challenge of those who want to know—if science is so credible—why it does not always provide the right answer (even in the natural sciences) and sometimes fails. If we can do this, we will not only understand what is distinctive about science, we will have the tools necessary to emulate its approach in other empirical fields too.

But there is another problem: we cannot pretend, these days, that the conclusions of science are going to be accepted just because they are rational and justified. Climate change skeptics insist that we need more evidence to prove global warming. Vaccine resisters maintain that there is a conspiracy to deny the truth about autism. What should we do about the problem of those who would simply reject the results of science? We may be tempted to dismiss these people as irrational, but we do so at our peril.

If we cannot provide a good account of why scientific explanations have a superior claim to believability, why should they accept them? It's not just that if we don't understand science we cannot cultivate it elsewhere; we cannot even defend science where it is working.

In short, I think that many of those who have written about science have mishandled the claim that science is special because they have not said enough about the failures of natural science, the potential for the social sciences, and the drawbacks of those fields that seek the mantle of science without embracing its ethos. This has led to failure to emulate science by those fields that would like to do so, and also to the irrational rejection of scientific conclusions by those who are motivated by their ideologies to think that their own views are just as good.

So what *is* distinctive about science? As I hope to show, what is most notable is the *scientific attitude* toward empirical evidence, which is as hard to define as it is crucial. To do science we must be willing to embrace a mindset that tells us that our prior beliefs, ideologies, and wishes do not matter in deciding what can pass the test of comparison with the evidence. This is no easy thing to mark off with a criterion of demarcation—neither does it pretend to be a proxy for "scientific method"—but I argue that it is essential to engaging in (and understanding) science. This is something that can be emulated by the social sciences and also helps to explain what is *not* scientific about intelligent design theory, the emptiness of denialism by those who wish to reject the evidence for climate change, and the folly of other conspiracy theories that purport to succeed where science is restrained by bona fide skepticism. At its heart, what is distinctive about science is that it *cares about evidence* and is *willing to change its theories on the basis of evidence*. It is not the subject or method of inquiry but the values and behavior of those who engage in it that makes science special. Yet this is a surprisingly complex thing to unravel, both in the history of the past successes of science and also in a program for how to make other fields more scientific in the future.

In the chapters to follow, I will show how the scientific attitude helps us with three main tasks: understanding science (chapters 1 through 6), defending science (chapters 7 and 8), and growing science (chapters 9 and 10). When done right, the philosophy of science is not just descriptive or even explanatory, but prescriptive. It helps to explain not just why science has been so successful in the past, but why evidential and experimental methods have so much potential value for other empirical fields in the realizing in that moment that it was possible to split the atom.[2] Inspiration in science, as in art, can come from a diversity of sources. Yet many hold that the results of science have a special claim to be believed because of the distinctive way that they can be rationally reconstructed *after the fact*. Thus it is not the way that scientific theories are *found* that gives them such great credibility, it is the process by which they can be *logically justified*.

Science textbooks provide a cleaned-up version of history. They give us the result of many centuries of scientific conflict and make us feel that the process inevitably led to our present enlightened understanding. Historians of science know this to be inaccurate, but it remains immensely popular, because of the great convenience that this account of "scientific method" provides in supporting not only the claim that the content of science is especially credible but also the idea that the process of scientific explanation can be emulated by other disciplines that wish to make their own empirical discoveries.[3]

Yet even if the classic five-step method proves too simple to get the job done, there are other ways that philosophers have sought to characterize what is distinctive about science, and some of these focus on methodology. Here it is important not to get confused. The claim that there is no universal one-size-fits-all "scientific method"—where we put in sensory observations at one end and get scientific knowledge at the other—does not contradict the idea that there could be some unique methodological feature of science. To say that there is no recipe or formula for producing scientific theories is a very different thing than to claim that scientists have *no methods whatsoever.* This is to say that even if most philosophers of science are willing to reject the idea of a simple "scientific method," many still think there is enormous benefit to analyzing the methodological differences between science and nonscience, in search of a way to justify the epistemic authority of those scientific theories that have already been discovered.

The Relevance of the Problem of Demarcation

One benefit of focusing on the methodology of science is that it purports to provide a way to demarcate between what is science and what is not. This is the so-called *problem of demarcation*, and it has been of enormous concern to the philosophy of science at least since the time of Karl Popper at the beginning of the twentieth century. In his essay "The Demise of the Demarcation

Problem," Larry Laudan claims that the problem of demarcation goes all the way back to Aristotle—who sought to differentiate between knowledge and opinion—and surfaced again in the era of Galileo and Newton—who pushed science into the modern era by using empirical methods to understand how nature worked. By the beginning of the nineteenth century, Laudan asserts, Auguste Comte and others began to hone in on the claim that what was distinctive about science was its "method," even if there was as yet no widespread agreement about what that actually was.[4] By the beginning of the twentieth century, philosophers were ready to sharpen this analysis and attempt to solve the problem of demarcation by creating a strict "criterion of demarcation" that could differentiate science from nonscience.

The Logical Positivists tried to do this on the basis of the allegedly special "meaning" of scientific statements. In contrast to other kinds of assertions, scientific statements were accepted as making a difference to our experience in the world, which meant that they must in some way be verifiable through sensory data. If scientists said that the planet Venus went through phases, we had better be able to see that through a telescope. Statements that could not do this (other than those used in logic, which is deductively valid and so already on firm ground) were branded "cognitively meaningless" because they were unverifiable; they were dismissed as nonsense and unworthy of our time, because there was no procedure for determining if they were right or wrong. If a statement about the world purported to be true, the positivists claimed, it must be verifiable by experience. If not, then instead of being scientific it was just "metaphysics" (which was the pejorative term they used to cover huge swaths of knowledge including religion, ethics, aesthetics, and the vast majority of philosophy). To pull off such a hard and fast distinction, though, the Logical Positivists needed to come up with a "verification criterion" by which one could sort meaningful from meaningless statements. And this, owing to technical reasons that ultimately came down to the problem that they couldn't get the sorting right, led to their undoing.[5]

The problem of demarcation was then taken up by perhaps its greatest champion, Karl Popper. Popper understood—even before Logical Positivism had formally failed—that there were problems with pursuing the verification of scientific statements. The positivists had based science on inductive inference, which undercut the idea that one could prove any empirical statement to be true. David Hume's famous problem of induction prevented the sort of logical certainty that the positivists coveted for scientific statements.[6] Even if they could not be proven, however, weren't scientific assertions nonetheless uniquely meaningful, given the fact that in principle they *could* be verified? Popper thought not, and regarded the pursuit of "meaning" as another mistake in the positivists' approach. What made science special, he felt, was not its meaning but its method. Popper thus set out, in the winter of 1919, to try to solve the problem of demarcation in another way—by renouncing both verification and meaning, focusing instead on what he called the "falsifiability" of scientific theories: the idea that they must be capable of being ruled out by some possible experience.

What concerned Popper was the difference between statements like those in astrology—that seemed compatible with any evidence—and those of science, that take some risk of being wrong. When an astrologer produces a personalized horoscope that says "You are sometimes insecure about your achievements and feel like an imposter" this can feel like a stunning insight into your inner-most thoughts, until you realize that the same horoscope is used for all clients. Contrast this with what happens in science. When a scientist makes a prediction, it comes with an understanding that if her theory is correct you will see what was predicted. And if you do *not* see that result, then the theory must be flawed.

Popper used this sort of contrast to think about the possible methodological difference between science and nonscience. He was searching for a way to forgo the impossibly high standard which said that scientific statements must always be *proven* by their evidence, but would still allow evidence to count. And then it hit him. If the Logical Positivists and others were searching for a way to differentiate science from nonscience—but were blocked from being able to say that scientific statements were *verifiable* because of Hume's problem of induction—why not instead follow the path of deductive certainty that was already enjoyed by logic?

Those who have studied formal logic know that the simplest and most famous deductively valid inference is *modus ponens*, which says "If A, then B. And A. Therefore B." No problem here. No need to check to see whether it "makes a difference to our experience." Deductive arguments are and always will be valid because the truth of the premises is sufficient to guarantee the truth of the conclusion; if the premises are true, the conclusion will be also. This is to say that the truth of the conclusion cannot contain any information that is not already contained in the premises. Consider the following valid argument:

If someone was born between 1945 and 1991, then they have Strontium-90 in their bones.

Adam was born in 1963.

Therefore, Adam has Strontium-90 in his bones.[7]

The problem with scientific statements, however, is that they don't seem to follow this form. For hundreds of years before Popper, they were accepted as being inductive, which meant that the reasoning looked more like "If A, then B. And B. Therefore A." For example:

If someone was born between 1945 and 1991, then they have Strontium-90 in their bones.

Eve has Strontium-90 in her bones.

Therefore, Eve was born between 1945 and 1991.

Obviously, this kind of argument is not deductively valid. The fact that Eve has Strontium-90 in her bones is no guarantee that she was born between 1945 and 1991. Eve might, for example, have grown up near a nuclear reactor in Pennsylvania in the late 1990s, where it was found that Strontium-90 was present as a result of environmental contamination. Here the form of the argument does not guarantee that if there the premises are true, the conclusion will be true. With inductive arguments, the conclusion contains information that goes beyond what is contained in the premises. This means that we will have to engage in some actual experience to see if the conclusion is true. But isn't this how we do science? Indeed, when we are engaged in reasoning about empirical matters, we often seek to go beyond our firsthand experiences and draw inferences about those situations that are similar to them. Even though our experience may be limited, we look for patterns within it and hope to be able to extrapolate them outward.

Suppose we are interested in a straightforward empirical issue such as the color of swans. We've seen a lot of swans in our life and they have all been white. We may therefore feel justified in making the assertion "All swans are white." Is this true? We've made our observations and have come up with a hypothesis, but now it is time to test it. So we make a prediction that from now on *every* swan we see will be white. Here is where it gets interesting. Suppose that this prediction turns out to be fulfilled. We may live our whole lives in North America, and, as it turns out, every single swan we ever see is white. Does this prove the truth of our assertion? No. It is still possible that someday if we go to Australia (or just open Google), we will see a black swan.

When we are trying to discover empirical truths about the world, we are hampered by the fact that our experience is always finite. No matter how long we live, we cannot possibly sample all of the swans who have lived or ever will live. So we can never be certain. If we wish to make blanket statements about the world—sometimes instantiated in scientific laws—we face the in principle worry that some future piece of evidence may come along to refute us. This is because the form of the argument that we are using here is inductive, and inductive inferences are not deductively valid. There is just no way to be certain that the rest of the world will conform to our limited experience.

Science nonetheless works pretty well. Although it may not guarantee the truth of our assertions, we are at least gathering evidence that is relevant to the warrant for our beliefs. And shouldn't this increase the likelihood that our general statements are true?[8] But why settle for this? Popper was bothered by the inductive form of inference used by positivists and others as the basis for science. But if that is its logical foundation, how can we possibly demarcate science from nonscience? To admit that "we could be wrong" doesn't sound like much of a distinction. Popper sought something stronger. He wanted a *logical* basis for the uniqueness of science.

Popper didn't have to look far. The inductive argument that we used above has a name—"affirming the consequent"—and it is a well-known fallacy of deductive logic. But there are other, better forms of argument, and one of the most powerful—*modus tollens*—is deductively valid. It works like this. "If A, then B. And not B. Therefore, not A."

If someone was born between 1945 and 1991, then they have Strontium-90 in their bones.

Gabriel does not have Strontium-90 in his bones.

Therefore, Gabriel was not born between 1945 and 1991.[9]

This was Popper's insight: this, he felt, was the logical basis for scientific inference. Just because science seeks to learn from empirical facts about the with intelligent design theory, which took a second swing at trying to get creationism into the public schools. For now the question is a philosophical one: could falsification identify what might be wrong with creation science? Some felt that it could, for just as with the earlier claims of astrology, it seemed that the main thesis of creation science—that God created the universe and all of the species within it—was compatible with *any* evidence. Didn't the discovery of 65-million-year-old dinosaur fossils conflict with the 6,000-year timeline in the Bible? Not really, the creation scientists contended, for surely an omnipotent God could have created the entire fossil record! I hope it is clear from our earlier consideration of the problems with astrology that this sort of tendency to explain away any contrary evidence is not a shining example of falsifiability. Whereas true science goes out on a limb to test its theories against experience, creation science refused to change its theory even when there was evidence against it. Add to this the fact that creation science had precious little to offer as positive evidence in its favor, and many were willing to dismiss it as nothing more than pseudoscience.[14]

The virtues of falsification are clear. If Popper had found a way to solve the problem of demarcation, philosophers and scientists now had a powerful tool for answering the question of what is special about science. They also had a mechanism for dismissing and criticizing those practices—such as astrology and creationism—that they did *not* want to accept as scientific; if they were not falsifiable, they were not scientific. An added benefit of Popper's approach was that he had found a way for a theory to be *scientific* without necessarily having to be *true*.[15] Why did this matter? In seeking a criterion of demarcation, it mattered a great deal to those who were versed in the history of science, and understood that some of the greatest scientific minds of the last few millennia had said things that later turned out to be false. It would be wrong to think that they weren't scientists. Even though Ptolemy's geocentric theory was later replaced by Copernicus's heliocentric one, this did not mean that Ptolemy wasn't a scientist. He had based his theory on empirical data and had pushed things forward as far as he could. What mattered is that his claims were *falsifiable*, not that they were later falsified.

It would be easy to imagine that Popper's new criterion of demarcation was also a vindication of the idea of "scientific method," but that would be far from true. In fact, Popper was one of the earliest and harshest critics of the idea that there *was* such a thing as "scientific method." In his most definitive statement on the subject, appropriately titled "On the Non-Existence of Scientific Method," Popper wrote "As a rule, I begin my lectures on Scientific Method by telling my students that scientific method does not exist."[16] Elsewhere, he writes:

> The belief that science proceeds from observation to theory is still so widely and so firmly believed that my denial of it is often met with incredulity. … But in fact the belief that we can start with pure observations alone, without anything in the nature of a theory, is absurd; as may be illustrated by the story of the man who dedicated his life to natural science, wrote down everything he could observe, and bequeathed his priceless collection of observations to the Royal Society to be used as inductive evidence. This story should show us that though beetles may profitably be collected, observations may not.[17]

It is important here to remember the distinction between saying that there is a "scientific method" and saying that there is some methodological difference—such as falsifiability—between science and nonscience. Although Popper is unequivocally rejecting the idea of "scientific method," he still believes that we can have a criterion of demarcation and even one that is methodological in nature.[18]

This opinion was not shared by some of Popper's critics, notably by one of his most famous, Thomas Kuhn, who felt that although Popper was correct to abandon the idea of scientific method,[19] one should probably also give up on the idea that there is *any* distinctive methodological difference between science and nonscience. Note that this does not necessarily mean that one is giving up on the idea that science is "special" or even that there is a way of distinguishing between science and nonscience. Kuhn was not yet ready to do this (though many of his later followers were); instead, he merely pointed out that the process by which scientists actually work has much more to do with nonevidential "subjective" factors in theory choice, such as scope, simplicity, fruitfulness, and the ability to fit a theory with one's other beliefs, and much less to do with any formal method. And surely this must have an impact on justification.

It is important to understand that Kuhn was not an opponent of science. He was not—although he has been blamed for it—one of those who later claimed that science was an "irrational" process, no better than any other way of knowing, nor did he believe that the social factors that sometimes influenced scientific theory choice undermined its claim to produce credible theories. Instead, Kuhn was at pains to make sure that we understood science for what it really was, feeling that even if we did so it would be no less wonderful. While Kuhn never took it upon himself to try to provide a criterion of demarcation, he did nonetheless feel himself to be a champion of science.[20]

What of Popper's theory? Despite its virtues, it was severely criticized—by Kuhn and others—as offering too simple a picture of scientific theory change, especially given the fact that most science did not work precisely as the heroic example of Einstein's prediction indicated. There are few such crucial tests, involving risky predictions and dramatic successes, in the history of science. Most of science actually grinds along fairly slowly, with tests on a much smaller scale and, tellingly, widespread reluctance to give up a workable hypothesis just because something has gone wrong.[21] Yes, evidence counts, and one cannot simply ignore data and insulate a theory from refutation. Yet many philosophers, embracing the Duhem–Quine thesis (which says that it is always easier to sacrifice a smaller supporting hypothesis or make an ad hoc modification than to give up a theory), were skeptical that science worked as Popper said it did. Even though Popper maintained that his theory dealt only with the *logical* justification of science, many felt that there was a growing credibility gap between the way that scientists actually worked and the way that philosophers justified their work, given the sorts of social factors that Kuhn had identified. As Kuhn demonstrated, we are occasionally given to engage in a scientific revolution, but it does not happen nearly often enough for this to be accepted as the basis for demarcating between science and nonscience.

The upshot of all this is that by the 1970s, there was fairly wide agreement among most philosophers of science not only that the classic five-step scientific method was a myth, but also that there was no genuine methodological distinction between science and nonscience. This had great importance for the idea that science was special. Can one defend the idea that science is distinctive without also believing in scientific method or at least some other criterion of demarcation? Many said no.

Once Kuhn had opened the door to examining the workaday details of how scientists did their business—through "puzzle solving" and the search for accommodation to the dominant paradigm through "normal science"—the critics seemed unstoppable. To Kuhn's horror (he did after all agree with Popper and other defenders of science that *evidence counts* and that the revolution from one scientific theory to another on the basis of this evidence is the hallmark of science), his work was often cited as support by those who no longer believed that science was special. Sociologists of science, relativists, postmodernists, social constructivists, and others took their turns attacking the idea that science was rational, that it had anything to do with the pursuit of truth, or indeed that scientific theories were anything more than a reflection of the political biases about race, class, and gender held by the scientists who produced them. To some, science became an ideology, and facts and evidence were no longer automatically accepted as providing credible grounds for theory choice even in the natural sciences.

Paul Feyerabend went so far as to claim that there was no method in science at all. This was a radical departure from merely giving up on the scientific method. Along for the ride went any claims about methodology (such as objectivity), a criterion of demarcation, and even the idea that scientific beliefs were privileged.[22] Many wondered whether philosophy had now given up on science all together.

This is not to suggest that all philosophers of science felt this way. There were many who followed the ideas of Logical Empiricism (the successor to Logical Positivism), which held sway contemporaneously with Popper's theory of falsification right through the Kuhnian revolution. Here the focus was on defending the special method of science—even picking up on the earlier positivist idea of a "unified science" whose method could be extended to the social sciences—not through falsifiability (or meaning), but through close examination of how one might build more credible and reliable theories even in the face of the problem of induction.[23] Even here, though, it was necessary to modulate the full-throated defense of science, and certain concessions had to be made.[24]

By 1983, Larry Laudan, one of the most prominent philosophers of science, was ready to pull the plug on the idea that one could have a criterion of demarcation. Laudan's work was not so radical as to suggest that science wasn't important. He was one of the post-Kuhnians who looked for a way to uphold the idea that science could make "progress," though certainly not toward "true" theories or in any way that suggested the hegemony of science over other ways of knowing. In his earlier-cited article "The Demise of the Demarcation Problem," Laudan argued that there was no possible solution to the problem of demarcation, largely on the grounds that if it could be solved it would have been by now. By the time Laudan entered the picture, it goes without saying that there is no scientific method, but even the idea of finding another way of distinguishing between science and nonscience now seems dead.

Note that this does not necessarily mean that there is no difference between science and nonscience. One could even believe (as I think Laudan does) that science is uniquely explanatory. It's just that we are not going to be able to find a workable device for demarcation. Even if we all agree in our bones what is science and what is not, we are not going to be able to create a hard and fast way to distinguish it. The technical reason for this, Laudan tells us, is that philosophers have not been able to come up with a set of necessary and sufficient conditions for science. And to him this seems to be an absolute requirement for fulfilling a criterion of demarcation.

> What will the formal structure of a demarcation criterion have to look like if it is to accomplish the tasks for which it is designed? Ideally, it would specify a set of individually necessary and jointly sufficient conditions for deciding whether an activity or set of statements is scientific or unscientific. As is well known, it has not proved easy to produce a set of necessary and sufficient conditions for science. Would something less ambitious do the job? It seems unlikely. Suppose, for instance, that someone offers us a characterization which purports to be a necessary (but not sufficient) condition for scientific status. Such a condition, if acceptable, would allow us to identify certain activities as decidedly unscientific, but it would not help "fix our beliefs," because it would not specify which systems actually were scientific. … For different reasons, merely sufficient conditions are equally inadequate. If we are told, "Satisfy these conditions and you will be scientific," we are left with no machinery for determining that a certain activity or statement is nonscientific. … Without conditions which are both necessary and sufficient, we are never in a position to say 'this is scientific: but that is nonscientific.[25]

What is the problem with giving only a necessary condition? It is too strict. By aiming to exclude all that is *not* science, we may also keep out some things that we would want to include. Suppose our necessary condition were that a science must be capable of performing controlled experiments. Doesn't that rule out geology? Astronomy? All of the social sciences? Suppose, on the other hand, that we abandon this and instead aim at providing only a sufficient condition for scientific investigation: for instance, that it must be concerned with seeking truth based on empirical evidence. Here the concern is that we may have been too inclusive. Haven't we now allowed as scientific the search for Bigfoot? By trying to include everything that *is* science, we may also let in those things that we would surely want to keep out.[26] Thus, Laudan tells us, to have an adequate criterion of demarcation

One imagines that if such scrupulously faithful adherence to the scholarship of the day had been allowed to hold sway in court the creationists would have been thrilled: yes, teach creationism as "bad" science, but teach it nonetheless.

Thus we see that failure to demarcate between science and nonscience can have real-world consequences outside philosophy. For one thing, the issue of teaching creationism in the public schools did not simply disappear in 1982, but instead has morphed and grown—partially as a result of philosophers' inability to defend what is special about science—into the current claim that "intelligent design (ID)" (which I have elsewhere referred to as "creationism in a cheap tuxedo") is now a full-fledged scientific theory that is ready for its debut in biology classrooms.[36] This was again put to test at trial in 2005 in *Kitzmiller v. Dover Area School District*, where another judge—in a stinging rebuke that was reminiscent of the *McLean* decision—found that intelligent design is "not science" and ordered the defendants to pay $1 million to the plaintiffs. This may give pause to future ID theorists, but sadly this story is still not over, as current "academic freedom" bills are pending in the state legislatures of Colorado, Missouri, Montana, and Oklahoma, modeled on a successful 2012 Tennessee law that defends the rights of "teachers who explore the 'scientific strengths and scientific weaknesses' of evolution and climate change."[37]

Figuring out how to demarcate between science and nonscience is no laughing matter. Being able to say, in public and in a comprehensible way, why science is special seems a particular duty for those philosophers of science who *believe in science*, but have not been able to articulate why. As climate change deniers begin to gear up, taking a page from the earlier battles of the creationists (and the tobacco lobby) in fighting scientific conclusions that they don't like through funding "junk science," then spreading it through public relations, isn't it time that we found a way to fight back?

Of late, this is precisely what has happened. In 2013, philosophers Massimo Pigliucci and Maarten Boudry published an anthology entitled *Philosophy of Pseudoscience: Reconsidering the Demarcation Problem*, in which they self-consciously seek to resurrect the problem of demarcation thirty years after Laudan's premature obituary. The papers are a treasure trove of the latest philosophical thinking on this issue, as the profession tries to steer its way out of the ditch where Laudan left it: where we believe that science is special, but can't quite say how. It is disappointing—but certainly understandable—that after all this time philosophers are a little unsure how to proceed. Perhaps resurrecting the traditional problem of demarcation is the answer. Or perhaps there is another way.

It is no small thing to dismiss the problem of demarcation, which has been the backbone of the philosophy of science since its founding. The attractiveness of using its structure and vocabulary as a way to understand and defend the distinctiveness of science is obvious. Perhaps this is why virtually all previous attempts to say what is special about science have involved trying to come up with some criterion of demarcation. But there are many pitfalls to resurrecting this approach.

In Pigliucci's essay "The Demarcation Problem: A (Belated) Response to Laudan," he rejects the "necessary and sufficient conditions" approach, preferring instead to rely on Ludwig Wittgenstein's concept of "family resemblance." Pigliucci thus claims to rescue the problem of demarcation from Laudan's "old-fashioned" approach (which may be conceived of as challenging Laudan's "meta-argument" over what is required to solve the problem of demarcation).[38] Instead Pigliucci's idea is to treat learning the difference between science and pseudoscience as a kind of "language game," where we come to learn the clusters of similarity and difference between different concepts by seeing how they are used. The goal here is to identify the various threads of relationship that do not fall neatly along the lines of necessary and sufficient conditions but nonetheless characterize what we mean when we say that some particular inquiry is scientific. Two of these threads—"empirical knowledge" and "theoretical understanding"—appear to do most of the work. As Pigliucci writes, "if there is anything we can all agree on about science, it is that science attempts to give an empirically based theoretical understanding of the world, so that a scientific theory has to have both empirical support and internal coherence and logic."[39] As a result, Pigliucci thinks that—among other things—we will have discovered a "Wittgensteinian family resemblance" for the concepts of science and pseudoscience that provides a viable demarcation criterion to "recover much (though not necessarily all) of the intuitive classification of sciences and pseudosciences generally accepted by practicing scientists and many philosophers of science."[40]

This account, however, seems quite nebulous as a criterion of demarcation. For one thing, what is its logical basis? At various points Pigliucci refers to the use of "fuzzy logic" (which relies on inferring degrees of membership for inclusion in a set) to help make his criterion more rigorous, but it remains unclear how this would work. As Pigliucci admits, "for this to actually work, one would have to develop quantitative metrics of the relevant variables. While such development is certainly possible, the details would hardly be uncontroversial."[41] To say the least: one imagines that the central concepts of empirical knowledge and theoretical understanding may be equally as difficult to describe and differentiate from their opposites as the concept of science itself. Has Pigliucci solved the problem of demarcation or merely pushed it back a step?

Others who have pursued a "post-Laudan" solution to the problem of demarcation have encountered a similarly rocky path. In the same volume, Sven Hansson pursues an extremely broad definition of science. Apparently wary of the implications of classifying disciplines like philosophy as nonscience, he instead expands the scope of science to mean something more like a "community of knowledge" and then proceeds to demarcate this from pseudoscience. For all of the alleged advantages of rescuing the humanities from the realm of nonscience, however, the cost is quite high, for now he cannot say that the problem with pseudoscience has anything to do with its bastardization of empirical standards (since at least part of what he now classifies as science is not empirical either).[42]

Maarten Boudry takes a similarly questionable step in saying that he thinks there are really two demarcation problems—the "territorial" and the "normative"—instead of just one. The former dispute he dismisses as sterile. It is just a matter of "turf" that concerns separating science from legitimate but nonempirical epistemic endeavors like history and philosophy. According to Boudry, the real dispute is between science and pseudoscience; this is where the normative issue arises, because this is where we face those disciplines that are just pretending to be sciences.[43] Yet this bifurcation of the problem of demarcation reveals a basic confusion between saying that a discipline is *non*scientific and saying that it is *un*scientific. Does Boudry mean to identify the "territorial" dispute as one between fields that are scientific and *non*scientific? If so, that is a highly idiosyncratic and misleading use of the term. The dispute that he appears to be searching for when he talks about the territorial problem of demarcation seems to be between science and what may be called "*un*science." Yet why would this be the proper alternative to the normative dispute? The more traditional interpretation of the demarcation debate—revealed in most scholarship in the field—is that between science and *non*science, or between science and *pseudo*science. These are the terms of art used by Popper, Laudan, and most everyone else.[44] Instead, Boudry seems to be creating a *new* demarcation problem, while saying nothing about why we should ignore the classic problem of demarcating between science and *non*science. But why does Boudry think that he can make a case for the normative battle between science and *pseudo*science, when he has not legitimately dispensed with the larger issue of science versus *non*science? The straw-man "territorial" distinction between science versus *un*science (history, philosophy, etc.) does not do the job.[45]

This struggle to explain whether the problem of demarcation should be between science versus *non*science—or between science versus *pseudo*science—may seem like a mere terminological dispute, but it is not. For if we are attempting to distinguish science from *all that is not science*, it may lead to a very different criterion of demarcation than if we seek to distinguish science *merely from its imposter*. The important point here is to recognize that, according to most scholars, the category of nonscience *includes* both those fields that are pseudo-scientific and those that are unscientific. An inquiry can be nonscientific either because it is merely pretending to be scientific (in which case it is *pseudo*scientific) or because it concerns matters where empirical data are not relevant (in which case it is *un*scientific).[46] (See figure 1.1.)

This failure to be specific about what one is differentiating science from, however, not only exists in the "post-Laudan" essays by Pigliucci, Hansson, and Boudry, but seems to reflect a deep equivocation in the literature that goes all the way back through Laudan to Karl Popper himself. Remember that in *The Logic of Scientific Discovery*, Popper says that he is demarcating science from math, logic, and "metaphysical speculation."[47] By the time he gets to *Conjectures and Refutations*, however, his target is pseudoscience. In Laudan's essay, he too slides back and forth between talking about nonscience and pseudoscience.[48]

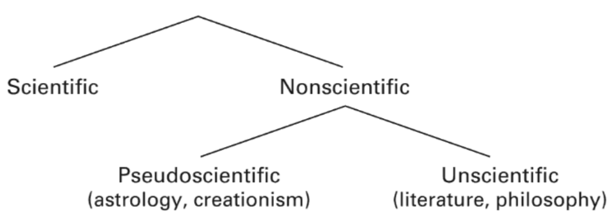

Figure 1.1

What difference does all of this make? It turns out to be crucial. Later we will revisit the issue of necessary and sufficient conditions and learn that the entire question of what is special about science may hang in the balance. We will see that the project of trying to solve the problem of demarcation is hamstrung unless we can specify precisely what it is that we are trying to define (science, nonscience, pseudoscience, or unscience) and, as we have seen, there has as yet been no definitive answer to that. My goal will be to provide a way to say what is distinctive about science without getting tripped up on the problem of providing both necessary and sufficient conditions—or trying to solve the problem of demarcation—because I do not think that these problems can be solved. Yet we still need a way to defend science.

First, however, let us deal with the problem of those who have misunderstood how science works.

carried out and (2) whether he can avoid becoming entangled in the need to rely on positive (confirming) instances. The Duhem–Quine thesis says that it is impossible to perform a "crucial test" in science, because every theory exists within a web of supporting assumptions. This means that even when there are falsifying instances for a theory, we can always sacrifice one of its auxiliary assumptions in order to save it. According to strict falsificationism, this must seem wrong. Popper himself, however, claims that he had already anticipated this objection and accommodated it within his theory.[1] According to the logic of falsification, when a theory is falsified we must abandon it. But in actual practice, Popper recognized, scientists are loath to give up a well-ensconced theory just because of one falsifying instance. Maybe they made a mistake; maybe there was something wrong with their apparatus. Popper admits as much when he writes "he who gives up his theory too easily in the face of apparent refutations will never discover the possibilities inherent in his theory."[2]

In explicating his falsificationist account, Popper nonetheless prefers to illustrate its virtues with stories of those larger-than-life dramatic instances when a theorist made a risky prediction that was later vindicated by the data. As we've seen, Einstein's general theory of relativity made a bold prediction about the bending of light in a strong gravitational field, which was confirmed during the total solar eclipse of 1919. If the prediction had been wrong, the theory would have been rejected. But, since it was right, the epistemic reward was tremendous.[3]

But most science does not work like this. In his book *Philosophy of Science: A Very Short Introduction*, Samir Okasha recounts the example of John Couch Adams and Urbain Le Verrier, who discovered the planet Neptune in 1846, when they were working (independently) within the Newtonian paradigm and noticed a slight perturbation in the orbit of the planet Uranus. Without Neptune, this would have been a falsifying instance for Newtonian theory, which held (following Kepler) that all planets should move in perfectly elliptical orbits, unless some other force was acting upon them. But, instead of abandoning Newton, Adams and Le Verrier instead sought out and found that other gravitational force.[4]

Some might complain that this is not even close to a falsifying instance, since the theorists were working well within the predictions that had been made by Newtonian theory. Indeed, Popper himself has sometimes cited this very example as just such an instance where scientists were wise not to give up on a theory too quickly. Yet contrast this with a similar challenge to Newtonian theory that had been around for over one hundred and fifty years by the time of the discovery of Neptune: the slight perihelion advance in the orbit of the planet Mercury.[5] How could this be explained? Astronomers tried various ad hoc solutions (sacrificing auxiliary assumptions) but none were successful. Finally, it was none other than Le Verrier himself who proposed that the slight perturbation in Mercury's orbit could be explained by the existence of an unseen planet—which he named Vulcan—between Mercury and the Sun. Although his tests to find it were unsuccessful, Le Verrier went to his grave in 1877 believing that Vulcan existed. Virtually all of the other astronomers at the time agreed with him in principle that—whether Vulcan existed or not—there must be some Newtonian explanation. Forty years later Einstein pulled this thread a little harder and the whole Newtonian sleeve fell off, for this was an instance in which the anomalous orbit was explained not by the strong gravitational force of another *planet*, but instead by the non-Newtonian idea that the *Sun*'s gravitational force could warp the space around it. When the orbit of Mercury fit those calculations, it ended up being a major mark in favor of the general theory of relativity. Yet this was not a prediction but a "retrodiction," which is to say that Einstein's theory was used to explain a past falsifying instance that Newtonian theorists *had been living with for over two hundred years*! How long is too long before one "gives up too easily" in the face of a falsified theory? Popper provides no rules for judging this. Falsification may work well with the logic of science, but it gives few guidelines for how scientists should actually choose between theories.

As Thomas Kuhn demonstrated, there are no easy answers to the question of when a falsifying instance should take down a well-believed theory and when it should merely lead us to keep searching for answers within the dominant paradigm. In Kuhn's work, we find a rich description of the way that scientists actually struggle with the day-to-day puzzles that present themselves in the course of "normal science," which is when we work to accommodate the predictions, errors, and what may seem like falsifying instances within the four corners of whatever well-accepted theory one may be working under at the time.[6] Of course, Kuhn also recognized that sometimes science does take a turn toward the dramatic. When anomalies pile up and scientists begin to have a hard time reconciling their paradigm with things it cannot explain, enough force builds up that a scientific revolution may occur, as the field shifts rapidly from one paradigm to another. But, as Kuhn tells us, this is often about more than mere lack of fit with the evidence; it can also encompass scope, simplicity, fruitfulness, and other "subjective" or "social" factors that Popper would be reluctant to account for within his logical account of falsification.

But there are other problems for Popper's theory, too, even if a hypothesis *passes* all of the rigorous tests we can throw at it. As Popper admits, even when a theory succeeds it cannot be accepted as true—or even approximately true—but must always inhabit the purgatory of having merely survived "so far."[7] As powerful as scientific testing is, we are always left at the end with a potentially infinite number of hypotheses that *could* fit the data, and an infinite supply of possible empirical evidence that *could* overthrow any theory. Scientific reasoning thus is forced to make peace with the fact that it will always be open ended, because the data will always be open ended too. Even if we are not "inductivists" about our method, we must admit that even when a theory has passed many tests there are always going to be more ahead.

Popper tried to solve this problem through his account of corroboration, in which he said that after a theory had survived many rigorous tests, it built up a kind of credibility such that we would be foolish to abandon it cavalierly; as Popper put it, some theories have "proven their mettle."[8] To some ears, however, this began to sound very much like the type of verification and confirmation that Popper insisted he had abandoned. Of course, we cannot say that a theory is *true* just because it has passed many tests. Popper admits as much. But the problem is that we also cannot say that a theory is even more *likely* to be true on the basis of these tests either. At some level, Popper seems to understand this danger (as well he should, for this is just the problem of induction again), but it is unclear how he proposes to deal with it.[9] Remember that induction undermines not just certainty, but also probability. If the sample of potential tests is infinite, then the sample we've chosen to test our theory against is infinitesimally small. Thus it cannot help a theory's verisimilitude that it is well corroborated. At various points, Popper says that falsification is a "purely logical affair."[10] But why then does he introduce a notion like corroboration? Perhaps this is a rare concession to the practical issues faced by scientists? Yet what to make then of Popper's claim that falsification is not concerned with practical matters?[11]

Philosophers of science will continue to fight over the soul of Karl Popper. Meanwhile, many are left to draw the inevitable conclusion that inductivist-type problems cannot help but accrete for surviving theories, even if one is a falsificationist. The idea that there are always more data that can overthrow a theory—along with the related idea that there are infinitely many potential theories to explain the evidence that we already have—is frequently dismissed as a philosopher's worry by many scientists, who continue to maintain that when a theory survives rigorous testing it is either true or more likely to be true. But in their bones I am not sure how they could possibly believe this, for scientists as well as philosophers of science know full well that the history of science is littered with the wreckage of such hubris.[12] Perhaps this is one instance in which some of those who wish to defend science may be tempted to pretend that science *can* prove a theory, even if they know it cannot. Sometimes, in the excitement of discovery, or the heat of criticism, it may seem expedient to maintain that one's theory is *true*—that certainty is possible. But I submit that those who wish to defend science have a special obligation to accept the peculiarities of science for what they are and not retreat behind half-truths and wishful thinking when called upon to articulate what is most special about it. Whether we believe that science is based on inductive reasoning or not—whether we believe in falsificationism or not—we must accept that (no matter how good the evidence) science *cannot* prove the truth of any empirical theory, nor can we say that it is even, technically speaking, more probably true.

Yet as we shall see, this does not mean that we have no grounds for believing a scientific theory.

"Just a Theory"

At this point it is important to confront a second misconception about science, which is that if a scientific claim falls short of "proof," "truth," or "verification" then it is "just a theory" and should not be believed. Sometimes this is read out as a claim that another theory is "just as likely to be true" or "could be true," while others maintain that any theoretical knowledge is simply inferior.

The first thing to understand is that there is a difference between a theory and a hypothesis. A hypothesis is in some ways a guess. It is not normally a wild guess; it is usually informed by some prior experience with the subject matter in question. Normally a hypothesis arises after someone has noticed the beginnings of a pattern in the data and said, "Huh, that's funny, I wonder if…" Then comes the prediction and the test. A hypothesis has perhaps been "back-tested" by our reason to see if it fits with the data we have encountered so far, but this is a far cry from the sort of scrutiny that will meet a hypothesis once it has become a theory.

A scientific theory must not only be firmly embedded in empirical evidence, it must also be capable of predictions that can be extrapolated into the wider world, so that we can see whether it survives a rigorous comparison with *new* evidence. The standards are high. Many theories meet their demise before they are even put forth by their proponents, as a result of meticulous self-testing in advance of peer review. Customarily a theory must also include an explanatory account for *why* the predictions are expected to work, so that there is a way to reason back from any falsification of a prediction to what might be wrong with the theory that produced it. (As we saw with the example of the perihelion advance of Mercury, it is also a plus if a theory can explain or retrodict any anomalies that scientists have been contending with in their previous theory.)

Here we should return to Karl Popper for a moment and give him some credit. Although he may not have been right in the details of his account of falsifiability—or his more general idea that one can demarcate science from nonscience on a logical basis—he did at least capture one essential element that explains how science works: our knowledge of the world grows by keeping close to the empirical evidence. We can love a theory only provisionally and must be willing to abandon it either when it is refuted or when the data favor a better one. This is to say that one of the most special things about science is that *evidence counts*.

The Nobel Prize–winning physicist Richard Feynman said it best:

> In general, we look for a new law by the following process: First we guess it.… Then we compute the consequences of the guess to see what…it would imply. And then we compare the computation results to nature, or we say compare to experiment or experience, compare it directly with observations to see if it works. If it disagrees with experiment, it's wrong. In that simple statement is the key to science. It doesn't make any difference how beautiful your guess is, it doesn't make any difference how smart you are, who made the guess, or what his name is. If it disagrees with experiment, it's wrong. That's all there is to it.[13]