

For my parents

Peter Godfrey-Smith is associate professor of philosophy and of history and philosophy of science at Stanford University. He is the author of *Complexity and the Function of Mind in Nature*.

The University of Chicago Press, Chicago 60637
The University of Chicago Press, Ltd., London
© 2003 by The University of Chicago
All rights reserved. Published 2003
Printed in the United States of America

12 11 10 09 08

4 5

ISBN: 0-226-30062-5 (cloth)
ISBN: 0-226-30063-3 (paper)

Library of Congress Cataloging-in-Publication Data

Godfrey-Smith, Peter.

Theory and reality : an introduction to the philosophy of science /
Peter Godfrey-Smith.

p. cm. — (Science and its conceptual foundations)

Includes bibliographical references and index.

ISBN 0-226-30062-5 (alk. paper) — ISBN 0-226-30063-3 (pbk. : alk.
paper)

I. Science—Philosophy. I. Title. II. Series.

Q175 .G596 2003

501—dc21

2002155305

Ⓢ The paper used in this publication meets the minimum requirements of the American National Standard for Information Sciences—Permanence of Paper for Printed Library Materials, ANSI Z39.48-1992.

in a way modeled on the understanding of the physical world achieved during the Scientific Revolution.

Further Reading

The topics in this chapter will be discussed in detail later, and references will be given then. Two other introductory books are worth mentioning, though. Hempel's *Philosophy of Natural Science* (1966) was for many years the standard introductory textbook in this area. It opens with the story of Semmelweis and is a clear and reasonable statement of mainstream twentieth-century empiricism. Alan Chalmers's *What Is This Thing Called Science?* (1999) is also very clear; it presents a different view from Hempel's and the one defended here.

For all the topics in this book, there are also reference works that readers may find helpful. Simon Blackburn's *Oxford Dictionary of Philosophy* is a remarkably useful book and is fun to browse through. The *Routledge Encyclopedia of Philosophy* is also of high quality. *The Blackwell Companion to the Philosophy of Science* has many short papers on key topics (though many of these papers are quite advanced). The *Stanford Online Encyclopedia of Philosophy* is still in progress but will be a very useful (and free) resource.

There are many good books on the Scientific Revolution, each with a different emphasis. Cohen, *The Birth of a New Physics* (1985), is a classic and very good on the physics. Henry, *The Scientific Revolution and the Origins of Modern Science* (1997), is both concise and thorough. It has an excellent chapter on mechanism and contains a large annotated bibliography. Schuster 1990 is also a useful quick summary, and Dear's *Revolutionizing the Sciences* (2001) is a concise and up-to-date book with a good reputation. But Toulmin and Goodfield's *Fabric of the Heavens* (1962), an old book recently reprinted, is my favorite. It focuses on the conceptual foundations underlying the development of scientific ideas. (It is the first of three books by Toulmin and Goodfield on the history of science; the second, *The Architecture of Matter* is also relevant here.)

Kuhn's *Copernican Revolution* (1957), is another classic, focused on the early stages, as the title suggests. Shapin's *Scientific Revolution* (1996), is not a good introduction to the Scientific Revolution but is a very interesting book anyway. There are several good books that focus on particular personalities. Koestler, *The Sleepwalkers* (1968), is fascinating on Kepler, and Sobel, *Galileo's Daughter* (1999), is also good on Galileo (and his daughter, a nun leading a tough life). The standard biography of the amazingly strange Isaac Newton, by Robert Westfall, comes in both long (1980) and short (1993) versions.

For a history of medicine, covering the whole world, see Porter, *The Greatest Benefit to Mankind* (1998).

2

Logic Plus Empiricism

2.1 The Empiricist Tradition

The first approach to science that we will examine is a revolutionary form of empiricism that appeared in the early part of the twentieth century, flourished for a time, was transformed and moderated under the pressure of objections, and then slowly became extinct. The earlier version of the view is called "logical positivism," and the later, moderate form is more usually called "logical empiricism." There is variation in terminology here; "logical empiricism" is sometimes used for the whole movement, early and late. Although we will be looking at fossils in this chapter, these remnants of the past are of great importance in understanding where we are now.

Before discussing logical positivism, it will be helpful to go even further back and say something about the empiricist tradition in general. In the first chapter I said that empiricism is often summarized with the claim that the only source of knowledge is experience. This idea goes back a long way, but the most famous stage of empiricist thought was in the seventeenth and eighteenth centuries, with the work of John Locke, George Berkeley, and David Hume. These "classical" forms of empiricism were based upon theories about the mind and how it works. Their view of the mind is often called "sensationalist." Sensations, like patches of color and sounds, appear in the mind and are all the mind has access to. The role of thought is to track and respond to patterns in these sensations. This view of the mind is not implied by the more basic empiricist idea that experience is the source of knowledge, but for many years such a view was common within empiricism.

Both during these classical discussions and more recently, a problem for empiricism has been a tendency to lapse into *skepticism*, the idea that we cannot know anything about the world. This problem has two aspects. One aspect we can call *external world skepticism*: how can we ever know anything about the real world that lies behind the flow of sensations? The

second aspect, made vivid by David Hume, is *inductive skepticism*: why do we have reason to think that the patterns in past experience will also hold in the future?

Empiricism has often shown a surprising willingness to throw in the towel on the issue of external world skepticism. (Hume threw in the towel on both kinds of skepticism, but that is unusual.) Many empiricists have been willing to say that they don't *care* about the possibility that there might be real things lying behind the flow of sensations. It's only the sensations that we have any dealings with. Maybe it makes no sense even to try to *think* about objects lying behind sensations. Perhaps our concept of the world is just a concept of a patterned collection of sensations. This view is sometimes called "phenomenalism." During the nineteenth century, phenomenalist views were quite popular within empiricism, and their oddity was treated with nonchalance. John Stuart Mill, an English philosopher and political theorist, once said that matter may be defined as "a Permanent Possibility of Sensation" (1865, 183). Ernst Mach, an Austrian physicist and philosopher, illustrated his phenomenalist view by drawing a picture of the world as it appeared through his left eye (see fig. 2.1; the shape in the lower right part of the image is his elegant mustache). All that exists is a collection of observer-relative sensory phenomena like these.

I hope phenomenalism looks strange to you, despite its eminent proponents. It is a strange idea. But empiricists have often found themselves backing into views like this. This is partly because they have often tended to think of the mind as *confined* behind a "veil of ideas" or sensations. The mind has no "access" to anything outside the veil. Many philosophers, including me, agree that this picture of the mind is a mistake. But it is not so easy to set up an empiricist view that entirely avoids the bad influence of this picture.

In discussions of the history of philosophy, it is common to talk of a showdown in the seventeenth and eighteenth centuries between "the rationalists" and "the empiricists." Rationalists like Descartes and Leibniz believed that pure reasoning can be a route to knowledge that does not depend on experience. Mathematics seemed to be a compelling example of this kind of knowledge. Empiricists like Locke and Hume insisted that experience is our only way of finding out what the world is like. In the late eighteenth century, a sophisticated intermediate position was developed by the German philosopher Immanuel Kant. Kant argued that all our thinking involves a subtle *interaction* between experience and preexisting mental structures that we use to *make sense* of experience. Key concepts like space, time, and causation cannot be derived from experience, because a person must *already* have these concepts in order to use experience to learn



Fig. 2.1

"The assertion, then, is correct that the world consists only of our sensations" (Mach 1897, 10).

about the world. Kant also held that mathematics gives us real knowledge of the world but does not require experience for its justification.

Empiricists must indeed avoid overly simple pictures of how experience affects belief. The mind does not passively receive the imprint of facts. The active and creative role of the mind must be recognized. The trick is to avoid this problem while still remaining true to basic empiricist principles.

As I said above, in the history of philosophy the term "rationalism" is often used for a view that opposes empiricism. In the more recent discussions of science that we are concerned with here, however, the term is generally not used in that way. (This can be a source of confusion; see the glossary.) The views called "rationalist" in the twentieth century were often forms of empiricism; the term was often used in a broad way, to indicate a confidence in the power of human reason.

So much for the long history of debate. Despite various problems, empiricism has been a very attractive set of ideas for many philosophers. Empiricism has often also had a particular kind of impact on discussions

outside of philosophy. Making a sweeping generalization, it is fair to say that the empiricist tradition has tended to be (1) pro-science, (2) worldly rather than religious, and (3) politically moderate or liberal (though these political labels can be hard to apply across times). David Hume, John Stuart Mill, and Bertrand Russell are examples of this tendency. Of the three elements of my generalization, religion is the one that has the most counter-examples. Berkeley was a bishop, for example, and Bas van Fraassen, one of the most influential living empiricist philosophers, is also religious. But on the whole it is fair to say that empiricist ideas have tended to be the allies of a practical, scientific, down-to-earth outlook on life. The logical positivists definitely fit this pattern.

2.2 The Vienna Circle

Logical positivism was a form of empiricism developed in Europe after World War I. The movement was established by a group of people who were scientifically oriented and who disliked much of what was happening in philosophy. This group has become known as the *Vienna Circle*.

The Vienna Circle was established by Moritz Schlick and Otto Neurath. It was based, as you might expect, in Vienna, Austria. From the early days through to the end, a central intellectual figure was Rudolf Carnap. Carnap seems to have been the kind of person whose presence inspired awe even in other highly successful philosophers.

Logical positivism was an extreme, swashbuckling form of empiricism. The term “positivism” derives from the nineteenth-century scientific philosophy of Auguste Comte. In the 1930s Carnap suggested that they change the name of their movement from “logical positivism” to “logical empiricism.” This change should not be taken to suggest that the later stages in the movement were “more empiricist” than the earlier stages. The opposite is true. In my discussion I will use the term “logical positivism” for the intense, earlier version of their ideas, and “logical empiricism” for the later, more moderate version. Although Carnap suggested the name change in the mid-1930s, the time during which logical positivist ideas changed most markedly was after World War II. I will spend some time in this section describing the unusual intellectual and historical context in which logical positivism developed. In particular, it is easier to understand logical positivism if we pay attention to what the logical positivists were *against*.

The logical positivists were inspired by developments in science in the early years of the twentieth century, especially the work of Einstein. They also thought that developments in logic, mathematics, and the philosophy of language had shown a way to put together a new kind of empiricist phi-

losophy that would settle, once and for all, the problems that philosophy had been concerned with. Some problems would be solved, and other problems would be rejected as meaningless. Logical positivist views about language were influenced by the early ideas of Ludwig Wittgenstein ([1922] 1988). Wittgenstein was an enigmatic, charismatic, and eccentric philosopher of logic and language who was not an empiricist at all. Some would say that the positivists adapted Wittgenstein’s ideas, others that they misinterpreted him.

Though they did admire some philosophers, the logical positivists were distressed with much of what had been going on in philosophy. In the years after Kant’s death in 1804, philosophy had seen the rise of a number of systems of thought that the logical positivists found pretentious, obscure, dogmatic, and politically harmful. One key villain was G. W. F. Hegel, who worked in the early nineteenth century and had a huge influence on nineteenth-century thought. Hegel was famous for his work on the relation between philosophy and history. He thought that human history as a whole was a process in which a “world spirit” gradually reached consciousness of itself. For Hegel, individuals are less important than the state as a whole, especially the role of the state in the grand march of historical progress. These ideas were taken to support strong forms of nationalism. Hegel’s was an “idealist” philosophy, since it held that reality is in some sense spiritual or mental. But this is not a view in which each person’s reality is made up in some way by that person’s ideas. Rather, a single reality *as a whole* is said to have a spiritual or rational character. This view is sometimes called “absolute idealism.”

Hegel’s influence bloomed and then receded in continental Europe. As it receded in continental Europe, in the later nineteenth century, it bloomed in England and America. Absolute idealism is a good example of what logical positivism was against. Sometimes the positivists would disparagingly dissect especially obscure passages from this literature. Hans Reichenbach (who was not part of the original Vienna Circle but who was a close ally) began his book *The Rise of Scientific Philosophy* (1951) with a quote from Hegel’s most famous work on philosophy and history: “Reason is substance, as well as infinite power, its own infinite material underlying all the natural and spiritual life; as also the infinite form, that which sets the material in motion.” Reichenbach lamented that a philosophy student, on first reading this passage, would usually think that it was *his* fault—the student’s fault—that he did not understand it. The student would then work away until it finally seemed obvious that Reason was substance, as well as infinite power. . . . For Reichenbach, it is entirely *Hegel’s* fault that the passage seems to make no sense. It seems to make no sense because whatever

factual meaning the claim might be intended to convey has been smothered with misused language.

People sometimes describe the history of this period as if it was a pitched battle between logical positivism and absolute idealism. That is not how things went. In the early twentieth century, there were many kinds of philosophy jostling and wrangling in Europe. There was a “back to Kant” movement going on (as there seems to be now; perhaps this will happen every hundred years). Another philosopher who came to seem an especially important rival to logical positivism was Martin Heidegger.

Earlier I gave a quick summary of Hegel’s ideas. It is much harder to do that for Heidegger. Heidegger is sometimes categorized as an existentialist. Perhaps he is the most famously difficult and obscure philosopher who has ever lived. I will borrow the summary reluctantly given by Thomas Sheehan in the entry for Heidegger in the *Routledge Encyclopedia of Philosophy* (1998): “He argues that mortality is our defining moment, that we are thrown into limited worlds of sense shaped by our being-towards-death, and that finite meaning is all the reality we get.” Simplifying even more, Heidegger held that we must understand our lives as based, first and foremost, upon practical coping with the world rather than knowledge of it. All our experience is affected by the awareness that we are traveling toward death. And the best thing we can do in this situation is stare it in the face and live an “authentic” life.

This picture of life might seem to make some sense (especially on a bad day). But Heidegger combined his descriptions of how it feels to live in the world with abstract metaphysical speculation; especially notorious are his discussions of the nature of “Nothing.” Heidegger also had one point in common with some (though not all) absolute idealists: his opposition to liberal democratic political ideas.

Heidegger was seen as a key rival by the logical positivists. Carnap gave humorous logical dissections of Heidegger’s discussions of Nothing in his lectures. Interestingly, recent work has shown that Carnap and Heidegger understood each other better than was once supposed (Friedman 2000).

Logical positivism was a plea for Enlightenment values, in opposition to mysticism, romanticism, and nationalism. The positivists championed reason over the obscure, the logical over the intuitive. The logical positivists were also internationalists, and they liked the idea of a universal and precise language that everyone could use to communicate clearly. Otto Neurath was the member of the group with the strongest political and social interests. He and various others in the group could be described as democratic socialists. They had a keen interest in some movements in art and architecture at the time, such as the Bauhaus movement. They saw this

work as assisting the development of a scientific, internationalist, and practical outlook on society (Galison 1990).

The Vienna Circle flourished from the mid-1920s to the mid-1930s. Logical positivist ideas were imported into England by A. J. Ayer in *Language, Truth, and Logic* (1936), a vivid and readable book that conveys the excitement of the time. Under the influence of logical positivism, and the philosophy of G. E. Moore and Bertrand Russell, English philosophy abandoned absolute idealism and returned to its traditional empiricist emphasis, an emphasis it has retained (more or less) ever since.

In continental Europe the story turned out differently. For we have now, remember, reached the 1930s. The development of logical positivism ran straight into the rise of Adolf Hitler.

Many of the Vienna Circle had socialist leanings, some were Jewish, and there were certainly no Nazis. So the logical positivists were persecuted by the Nazis, to varying degrees. The Nazis encouraged and made use of pro-German, anti-liberal philosophers, who also tended to be obscure and mystical. In contrast to the logical positivists, Martin Heidegger joined the Nazi party in 1933 and remained a member throughout the war.

Many logical positivists fled Europe, especially to the United States. Schlick, unfortunately, did not. He was murdered by a deranged former student in 1936. The logical positivists who did make it to the United States were responsible for a great flowering of American philosophy in the years after World War II. These include Rudolf Carnap, Hans Reichenbach, Carl Hempel, and Herbert Feigl. In the United States the strident voice of logical positivists was moderated. Partly this was because of criticisms of their ideas—criticisms from the side of those who shared their general outlook. But the moderation was no doubt partly due to the different intellectual and political climate in the United States. Austria and Germany in the 1930s had been an unusually intense environment for doing philosophy.

2.3 Central Ideas of Logical Positivism

Logical positivist views about science and knowledge were based on a general theory of language; we need to start here, before moving to the views about science. This theory of language featured two main ideas, the *analytic-synthetic distinction* and the *verifiability theory of meaning*.

The analytic-synthetic distinction will probably strike you as bland and obvious, at least at first. Some sentences are true or false simply in virtue of their meaning, regardless of how the world happens to be; these are analytic. A synthetic sentence is true or false in virtue of both the meaning of the sentence *and* how the world actually is. “All bachelors are unmarried”

is the standard example of an analytically true sentence. “All bachelors are bald” is an example of a synthetic sentence, in this case a false one. Analytic truths are, in a sense, empty truths, with no factual content. Their truth has a kind of necessity, but only because they are empty.

This distinction had been around, in various forms, since at least the eighteenth century. The terminology “analytic-synthetic” was introduced by Kant. Although the distinction itself looks uncontroversial, it can be made to do real philosophical work. Here is one crucial piece of work the logical positivists saw for it: they claimed that all of mathematics and logic is analytic. This made it possible for them to deal with mathematical knowledge within an empiricist framework. For logical positivism, mathematical propositions do not describe the world; they merely record our conventional decision to use symbols in a particular way. Synthetic claims about the world can be expressed using mathematical language, such as when it is claimed that there are nine planets in the solar system. But proofs and investigations within mathematics itself are analytic. This might seem strange because some proofs in mathematics are very surprising. The logical positivists insisted that once we break down such a proof into small steps, each step will be trivial and unsurprising.

Earlier philosophers in the rationalist tradition had claimed that some things can be known a priori; this means known *independently of experience*. Logical positivism held that the only things that seem to be knowable a priori are analytic and hence empty of factual content.

A remarkable episode in the history of science is important here. For many centuries, the geometry of the ancient Greek mathematician Euclid was regarded as a shining example of real and certain knowledge. Immanuel Kant, inspired by the immensely successful application of Euclidean geometry to nature in Newtonian physics, even claimed that Euclid’s geometry (along with the rest of mathematics) is both synthetic and knowable a priori. In the nineteenth century, mathematicians did work out alternative geometrical systems to Euclid’s, but they did so as a mathematical exercise, not as an attempt to describe how lines, angles, and shapes work in the actual world. Early in the twentieth century, however, Einstein’s revolutionary work in physics showed that a non-Euclidean geometry is true of our world. The logical positivists were enormously impressed by this development, and it guided their analysis of mathematical knowledge. The positivists insisted that pure mathematics is analytic, and they broke geometry into two parts. One part is purely mathematical, analytic, and says nothing about the world. It merely describes possible geometrical systems. The other part of geometry is a set of synthetic claims about which geometrical system applies to our world.

I turn now to the other main idea in the logical positivist theory of language, the *verifiability theory of meaning*. This theory applies only to sentences that are not analytic, and it involves a specific kind of “meaning,” the kind involved when someone is trying to say something about the world. Here is how the theory was often put: *the meaning of a sentence consists in its method of verification*. That formulation might sound strange (it always has to me). Here is a formulation that sounds more natural: knowing the meaning of a sentence is knowing how to verify it. And here is a key application of the principle: if a sentence has no possible method of verification, it has no meaning.

By “verification” here, the positivists meant verification *by means of observation*. Observation in all these discussions is construed broadly, to include all kinds of sensory experience. And “verifiability” is not the best word for what they meant. A better word would be “testability.” This is because testing is an attempt to work out whether something is true or false, and that is what the positivists had in mind. The term “verifiable” generally only applies when you are able to show that something is true. It would have been better to call the theory “the testability theory of meaning.” Sometimes the logical positivists did use that phrase, but the more standard name is “verifiability theory,” or just “verificationism.”

Verificationism is a strong empiricist principle; experience is the only source of meaning, as well as the only source of knowledge. Note that verifiability here refers to verifiability in *principle*, not in practice. There was some dispute about which hard-to-verify claims are really verifiable in principle. It is also important that *conclusive* verification or testing was not required. There just had to be the possibility of finding observational evidence that would count for or against the proposition in question.

In the early days of logical positivism, the idea was that in principle one could *translate* all sentences with factual meaning into sentences that referred only to sensations and the patterns connecting them. This program of translation was fairly quickly abandoned as too extreme. But the verifiability theory was retained after the program of translation had been dropped.

The verifiability principle was used by the logical positivists as a philosophical weapon. Scientific discussion, and most everyday discussion, consists of verifiable and hence meaningful claims. Some other parts of language are clearly not intended to have factual meaning, so they fail the verifiability test but do so in a harmless way. Included are poetic language, expressions of emotion, and so on. But there are also parts of language that are *supposed* to have factual meaning—are supposed to say something about the world—but which *fail* to do so. For the logical positivists, this includes most traditional philosophy, much of ethics, and theology as well!

This analysis of language provided the framework for the logical positivist philosophy of science. Science itself was seen as just a more complex and sophisticated version of the sort of thinking, reasoning, and problem-solving that we find in everyday life—and completely *unlike* the meaningless blather of traditional philosophy.

So let us now look at the logical positivists' picture of science and of the role of philosophy in a scientific worldview. Next we should turn to another distinction they made, between "observational" language and "theoretical" language. There was uncertainty about how exactly to set this distinction up. Usually it was seen as a distinction applied to individual terms. "Red" is in the observational part of language, and "electron" is in the theoretical part. There was also a related distinction at the level of sentences. "The rod is glowing red" is observational, while "Helium atoms each contain two electrons" is theoretical. A more important question was where to draw the line. Schlick thought that only terms referring to sensations were observational; everything else was theoretical. Here Schlick stayed close to traditional empiricism. Neurath thought this was a mistake and argued that terms referring to many ordinary physical objects are in the observational part of language. For Neurath, scientific testing must not be understood in a way that makes it private to the individual. Only observation statements about physical objects can be the basis of public or "inter-subjective" testing.

The issue became a constant topic of discussion. In time, Carnap came to think that there are lots of acceptable ways of marking out a distinction between the observational and theoretical parts of language; one could use whichever is convenient for the purposes at hand. This was the start of a more general move that Carnap made toward a view based on the "tolerance" of alternative linguistic frameworks.

We now need to look at logical positivist views about logic. For logical positivism, *logic is the main tool for philosophy*, including philosophical discussion of science. In fact, just about the only useful thing that philosophers can do is give logical analyses of how language, mathematics, and science work.

Here we should distinguish two kinds of logic (this discussion will be continued in chapter 3). Logic in general is the attempt to give an abstract theory of what makes some arguments compelling and reliable. Deductive logic is the most familiar kind of logic, and it describes patterns of argument that transmit truth with certainty. These are arguments with the feature that if the premises of the argument are true, the conclusion must be true. Impressive developments in deductive logic had been under way since the late nineteenth century and were still going on at the time of the Vienna Circle.

The logical positivists also believed in a second kind of logic, a kind that was (and is) much more controversial. This is *inductive* logic. Inductive logic was supposed to be a theory of arguments that provide support for their conclusions but do not give the kind of guarantee found in deductive logic.

From the logical positivist point of view, developing an inductive logic was of great importance. Hardly any of the arguments and evidence that we confront in everyday life and science carry the kind of guarantees found in deductive logic. Even the best kind of evidence we can find for a scientific theory is not completely decisive. There is always the possibility of error, but that does not stop some claims in science from being supported by evidence. The logical positivists accepted and embraced the fact that error is always possible. Although some critics have misinterpreted them on this point, the logical positivists did *not* think that science ever reaches absolute certainty.

The logical positivists saw the task of logically analyzing science as sharply distinct from any attempt to understand science in terms of its history or psychology. Those are empirical disciplines, and they involve a different set of questions from those of philosophy.

A terminology standardly used to express the separations between different approaches here was introduced by Hans Reichenbach. Reichenbach distinguished between the "context of discovery" and the "context of justification." That terminology is not helpful, because it suggests that the distinction has to do with "before and after." It might seem that the point being made is that discovery comes first and justification comes afterward. That is not the point being made (though the logical positivists were not completely clear on this). The key distinction is between the study of the logical structure of science and the study of historical and psychological aspects of science.

So logical positivism tended to dismiss the relevance of fields like history and psychology to the philosophy of science. In time this came to be regarded as a big mistake.

Let us put all these ideas together and look at the picture of science that results. Logical positivism was a revolutionary, uncompromising version of empiricism, based largely on a theory of language. The aim of science—and the aim of everyday thought and problem-solving as well—is to track and anticipate patterns in experience. As Schlick once put it, "what every scientist seeks, and seeks alone, are . . . the rules which govern the connection of experiences, and by which alone they can be predicted" (1932-33, 44). We can make rational predictions about future experiences by attending to patterns in past experience, but we never get a guarantee. We could always be wrong. There is no alternative route to knowledge besides experience;

when traditional philosophy has tried to find such a route, it has lapsed into meaninglessness.

The interpretation of logical positivism I have just given is a standard one. There is controversy about how to interpret the aims and doctrines of the movement, however. Some recent writers have argued that there is less of a link between logical positivism and traditional empiricism than the standard interpretation claims (Friedman 1999). But in the sense of empiricism used in this book, there is definitely a strong link. We see that in the Schlick quote given in the previous paragraph.

During the early twentieth century, there were various other strong versions of empiricism being developed as well. One was *operationalism*, which was developed by a physicist, Percy Bridgman (1927). Operationalism held that scientists should use language in such a way that all theoretical terms are tied closely to direct observational tests. This is akin to logical positivism, but it was expressed more as a proposed *tightening up* of scientific language (motivated especially by the lessons of Einstein's theory of relativity) than as an analysis of how all science already works.

In the latter part of the twentieth century, an image of the logical positivists developed in which they were seen as stodgy, conservative, unimaginative science-worshippers. Their strongly pro-science stance has even been seen as antidemocratic, or aligned with repressive political ideas. This is very unfair, given their actual political interests and activities. Later we will see how ideas about the relation between science and politics changed through the twentieth century in a way that made this interpretation possible. The accusation of stodginess is another matter; the logical positivists' writings were often extremely dry and technical. Still, even the driest of their ideas were part of a remarkable program that aimed at a massive, transdisciplinary, intellectual housecleaning. And their version of empiricism was organized around an ideal of intellectual flexibility as a mark of science and rationality. We see this in a famous metaphor used by Neurath (who exemplifies these themes especially well). Neurath said that in our attempts to learn about the world and improve our ideas, we are "like sailors who have to rebuild their ship on the open sea." The sailors replace pieces of their ship plank by plank, in a way that eventually results in major changes but which is constrained by the need to keep the ship afloat during the process.

2.4 Problems and Changes

Logical positivist ideas were always in a state of flux, and they were subject to many challenges. One set of problems was internal to the program. For example, there was considerable difficulty in getting a good formulation of

the verifiability principle. It turned out to be hard to formulate the principle in a way that would exclude all the obscure traditional philosophy but include all of science. Some of these problems were almost comically simple. For example, if "Metals expand when heated" is testable, then "Metals expand when heated and the Absolute Spirit is perfect" is also testable. If we could empirically show the first part of the claim to be false, then the whole claim would be shown false, because of the logic of statements containing "and." (If *A* is false then *A*&*B* must be false too.) Patching this hole led to new problems elsewhere; the whole project was quite frustrating (Hempel 1965, chap. 4). The attempt to develop an inductive logic also ran into serious trouble. That topic will be covered in the next chapter.

Other criticisms were directed not at the details but at the most basic ideas of the movement. The criticism that I will focus on here is one of these, and its most famous presentation is in a paper sometimes regarded as the most important in all of twentieth-century philosophy: W. V. Quine's "Two Dogmas of Empiricism" (1953).

Quine argued for a *holistic* theory of testing, and he used this to motivate a holistic theory of meaning as well. In describing the view, first I should say something about holism in general. Many areas of philosophy contain views that are described using the term "holism." A holist argues that you cannot understand a particular thing without looking at its place in a larger whole. In the case we are concerned with here, holism about testing says that we cannot test a single hypothesis or sentence in isolation. Instead, we can only test complex networks of claims and assumptions. This is because only a complex network of claims and assumptions makes definite predictions about what we should observe.

Let us look more closely at the idea that individual claims about the world cannot be tested in isolation. The idea is that in order to test one claim, you need to make assumptions about many other things. Often these will be assumptions about measuring instruments, the circumstances of observation, the reliability of records and of other observers, and so on. So whenever you think of yourself as testing a single idea, what you are really testing is a long, complicated *conjunction* of statements; it is the whole conjunction that gives you a definite prediction. If a test has an unexpected result, then something in that conjunction is false, but the failure of the test itself does not tell you *where* the error is.

For example, suppose you want to test the hypothesis that high air pressure is associated with fair, stable weather. You make a series of observations, and what you seem to find is that high pressure is instead associated with unstable weather. It is natural to suspect that your original hypothesis was wrong, but there are other possibilities as well. It might be that your

barometer does not give reliable measurements of air pressure. There might also be something wrong with the observations made (by you or others) of the weather conditions themselves. The unexpected observations are telling you that *something* is wrong, but the problem might lie with one of your background assumptions, not with the hypothesis you were trying to test.

Some parts of this argument are convincing. It is true that only a network of claims and assumptions, not a single hypothesis alone, tells us what we should expect to observe. The failure of a prediction will always have a range of possible explanations. In that sense, testing is indeed holistic. But this leaves open the possibility that we might often have good reasons to lay the blame for a failed prediction at one place rather than another. In practice, science seems to have some effective ways of working out where to lay the blame. Giving a philosophical theory of these decisions is a difficult task, but the mere fact that failed predictions always have a range of possible explanations does not settle the holism debate.

Holist arguments had a huge effect on the philosophy of science in the middle of the twentieth century. Quine, who sprinkled his writings with deft analogies and dry humor, argued that mainstream empiricism had been committed to a badly simplistic view of testing. We must accept, as Quine said in a famous metaphor, that our theories “face the tribunal of sense-experience . . . as a corporate body” (1953, 41). Logical positivism must be replaced with a holistic version of empiricism.

But there is a puzzle here. The logical positivists *already accepted* that testing is holistic in the sense described above. Here is Herbert Feigl, writing in 1943: “No scientific assumption is testable in complete isolation. Only whole complexes of inter-related hypotheses can be put to the test” (1943, 16). Carnap had been saying the same thing (1937, 318). We can even find statements like this in Ayer’s *Language, Truth, and Logic* (1936).

Quine did recognize Pierre Duhem, a much earlier French physicist and philosopher, as someone who had argued for holism about testing. (Holism about testing is often called “the Duhem-Quine thesis.”) But how could it be argued that logical positivists had dogmatically missed this important fact, when they repeatedly expressed it in print? Regardless of this, many philosophers agreed with Quine that logical positivism had made a bad mistake about testing in science.

Though the history of the issue is strange, it might be fair to say this: although the logical positivists officially accepted a holistic view about testing, they did not appreciate the significance of the point. The verifiability principle *seems* to suggest that you can test sentences one at a time. It seems to attach a set of observable outcomes of tests to each sentence in isolation.

Strictly, the positivists generally held that these observations are only associated with a specific hypothesis *against a background of other assumptions*. But then it seems questionable to associate the test results solely with the hypothesis itself. Quine, in contrast, made the consequences of holism about testing very clear. He also drew conclusions about language and meaning; given the link between testing and meaning asserted by logical positivism, holism about testing leads to holism about meaning. And holism about meaning causes problems for many logical positivist ideas.

The version of holism that Quine defended in “Two Dogmas” was an extreme one. It included an attack on the one idea in the previous section that you might have thought was completely safe: the analytic-synthetic distinction. Quine argued that this distinction *does not exist*; this is another unjustified “dogma” of empiricism.

Here again, some of Quine’s arguments were directed at a version of the analytic-synthetic distinction that the logical positivists no longer held. Quine said that the idea of analyticity was intended to treat some claims as *immune to revision*, and he argued that in fact no statement is immune to revision. But Carnap had already decided that analytic statements can be revised, though they are revised in a special way. A person or community can decide to drop one whole linguistic and logical framework and adopt another. Against the background provided by a given linguistic and logical framework, some statements will be analytic and hence not susceptible to empirical test. But we can always change frameworks. By the time that Quine was writing, Carnap’s philosophy was based on a distinction between changes made *within* a linguistic and logical framework, and changes *between* these frameworks.

In another (more convincing) part of his paper, Quine argued that there is no way to make *scientific sense* of a sharp analytic-synthetic distinction. He connected this point to his holism about testing. For Quine, all our ideas and hypotheses form a single “web of belief,” which has contact with experience only as whole. An unexpected observation can prompt us to make a great variety of possible changes to the web. Even sentences that might look analytic can be revised in response to experience in some circumstances. Quine noted that strange results in quantum physics had suggested to some that revisions in logic might be needed.

In this discussion of problems for logical positivism, I have included some discussions that started early and some that took place after World War II, when the movement had begun its U.S.-based transformation. Let us now look at some central ideas of logical empiricism, the later, less aggressive stage of the movement.

2.5 Logical Empiricism

Let's see how things looked in the years after World War II. Schlick is dead, and other remnants of the Vienna Circle are safely housed in American universities—Carnap at Chicago, Hempel at Pittsburgh and then Princeton, Reichenbach at UCLA (via Turkey), Feigl at Minnesota. Many of the same people are involved, but the work is different. The revolutionary attempt to destroy traditional philosophy has been replaced by a program of careful logical analysis of language and science. Discussion of the contributions that could be made by the scientific worldview to a democratic socialist future have been dropped or greatly muted. (Despite this, the FBI collected a file on Carnap as a possible Communist sympathizer.)

As before, ideas about language guided logical empiricist ideas about science. The analytic-synthetic distinction had not been rejected, but it was regarded as questionable. The logical empiricists felt the pressure of Quine's arguments. The verifiability theory, which had been so scythe-like in its early forms, was replaced by a *holistic empiricist theory of meaning*. Theories were seen as abstract structures that connect many hypotheses together. These structures are connected, as wholes, to the observable realm, but each *bit* of a theory—each claim or hypothesis or concept—does not have some specific set of observations associated with it. A theoretical term (like “electron” or “gene”) derives its meaning from its place in the whole structure and from the structure's connection to the realm of observation.

Late in the logical empiricist era, in 1970, Herbert Feigl gave a pictorial representation of what he called “the orthodox view” of theories (see fig. 2.2). A network of theoretical hypotheses (“postulates”) is connected by stages to what Feigl calls the “soil” of experience. This anchoring is the source of the network's meaning. Feigl used this picture to describe a single scientific theory. For the more extreme holism of Quine, a person's *total* set of beliefs form a *single* network.

The logical positivist distinction between observational and theoretical parts of language was kept roughly intact. But the idea that observational language describes private sensations had been dropped. The observational base of science was seen as made up of descriptions of observable physical objects (though Carnap thought it might occasionally be useful to work with a language referring to sensations).

Logical positivist views about the role of logic in philosophy and about the sharp separation between the logic of science and the historical-psychological side of science were basically unchanged. A good example of the kind of work done by logical empiricists is provided by their work on




Fig. 2.2

Feigl's picture of the logical empiricist view of theories
(From Feigl 1970; reproduced courtesy
of University of Minnesota Press)

explanation in science (see especially Hempel and Oppenheim 1948; Hempel 1965). For Hempel, to explain something is to show how to *infer* it using a logical argument, where the premises of the argument include at least one statement of a natural law (see chapter 13 below). This illustrates the idea, common to logical positivism and logical empiricism, that logic is the main tool of philosophy of science.

We saw that logical positivism held that the sole aim of science is to track patterns in experience. For logical positivism, when a scientist seems to be trying to describe unobservable structures in the world that give rise to what we see, the scientist must instead be seen as describing the *observable* world in a special, abstract way. Scientific language is only meaningful insofar as it picks out patterns in the flow of experience. Now, does logical *empiricism* make the same claim? Does logical empiricism claim that scientific language ultimately only describes patterns in observables?

The answer is that logical empiricists agonized over this. In their hearts their answer was *yes*, but this answer seemed to get harder and harder to defend. Carl Hempel wrote a paper in 1958 called “The Theoretician's Dilemma,” which was the height of logical empiricist agony over the issue. As a fairly traditional empiricist, Hempel was attracted to the idea that the

only possible role for those parts of language that seem to refer to unobservable entities is to help us pick out patterns in the observable realm. And if the parts of theories that appear to posit unobservable things are really any good, this “goodness” has to show up in advantages the theory has in its handling of observables. So there is no justification for seeing these parts of scientific language as describing real objects lying beyond experience. But Hempel and the logical empiricists found themselves forced to concede that this view does not make much sense of actual scientific work. When scientists use terms like “electron” or “gene,” they act as if they are doing more than tracking complex patterns in the observable realm. But the idea that the logical empiricists were being pushed toward—the idea that scientific theories are aimed at describing unobservable real structures—was hard to put on the table and defend. Empiricist philosophy of language seemed implacably opposed to it.

Empiricists were familiar with bad versions of the idea that behind the ordinary world of observables there is a special and superior realm, pure and perfect. This “layered” view of reality seemed to empiricists a source of endless trouble, right from the time of the ancient Greek philosopher Plato, who distinguished the illusory, unstable world of “appearances” from the more perfect and real world of “forms.” Empiricists have rightly been determined to avoid this kind of picture. But much of science does appear to be a process in which people hypothesize hidden structures that give rise to observable phenomena. These hidden structures are not “pure and perfect” or “more real” than the observable parts of the world, but they do lie behind or beneath observable phenomena. Of course, unobservable structures posited by a theory at one time might well turn out to be observable at a later time. In science, there is no telling what kinds of new access to the hidden parts of the world we might eventually achieve. But still, much of science does seem to proceed by positing entities that are, at the time of the research in question, truly hidden. For the traditional empiricist philosopher, understanding scientific theorizing in a way that posits a layer of observable phenomena and a layer of hidden structure responsible for the phenomena takes us *far too close* to bad old philosophical views like Plato’s. We are too close for comfort, so we must give a different kind of description of how science works.

The result is the traditional empiricist insistence that, ultimately, the only thing scientific language can do is describe patterns in the observable realm. In the first published paper that introduced logical positivism, Carnap, Hahn, and Neurath said: “In science there are no ‘depths’; there is surface everywhere” ([1929] 1973, 306). This is a vivid expression of the empiricist aversion to a view in which the aim of theorizing is to describe hidden

levels of structure. Science uses unusual theoretical concepts (which *look* initially like attempts to refer to hidden things) as a way of discovering and describing subtle patterns in the observable realm. So the logical positivists and the logical empiricists talked constantly about *prediction* as the goal of science. Prediction was a substitute for the more obvious-looking—but ultimately forbidden—goal of describing the real hidden structure of the world.

Twentieth-century empiricism made an important mistake here. We can make sense of science only by treating much of it as an attempt to describe hidden structures that give rise to observable phenomena. This is a version of *scientific realism*, an idea that will be discussed later in this book. In science there *are* depths. There is not a simple and fixed distinction between two “layers” in nature—the empiricists were right to distrust this idea. Instead there are *many* layers, or rather a *continuum* between structures that are more accessible to us and structures that are less accessible. Genes are hidden from us in some ways, but not as hidden as electrons, which in turn are not as hidden as quarks. Although there are “depths” in science, what is deep at one time can come to the surface at later times, and there may be lots of ways of interacting with what is presently deep.

2.6 On the Fall of Logical Empiricism

Logical empiricist ideas dominated much American philosophy, and they were very influential elsewhere in the English-speaking world and in some parts of Europe, in the middle of the twentieth century. But by the mid-1960s the view was definitely under threat; and by the middle or late 1970s, logical empiricism was near to extinction. The fall of logical empiricism was due to several factors, all of which I have either introduced in this chapter or will discuss in later chapters. One is the breakdown of the view of language that formed the basis of many logical positivist and logical empiricist ideas. Another is pressure from holist arguments. A third is the frustrating history of attempts to develop an inductive logic (chapter 3). A fourth is the development of a new role for fields like history and psychology in the philosophy of science (chapters 5–7). And eventually there was pressure from scientific realism. But this was only possible after logical empiricism had begun to decline.

Further Reading

For much more on the empiricist tradition in general, see Garrett and Barbanell, *Encyclopedia of Empiricism* (1997).

Schlick's "Positivism and Realism" (1932–33) and Feigl's "Logical Empiricism" (1943) are good statements of logical positivism by original members of the Vienna Circle. (Feigl uses the term "logical empiricism," but his paper describes a fairly strong, undiluted version of the view.) Ayer's *Language, Truth, and Logic* (1936) is readable, vivid, and exciting. Some see it as a distortion of logical positivist ideas.

The *Routledge Encyclopedia of Philosophy* (1998) has an interesting collection of articles, especially in the light of new debates about the history of logical positivism. The article on logical positivism is by Friedman and reflects his somewhat unorthodox reading (de-emphasizing the empiricist tradition). Stadler's entry on the Vienna Circle gives a more traditional view. See also Creath's entry on Carnap. On all these issues, see also the essays in Giere and Richardson 1997.

Peter Galison's "Aufbau/Bauhaus" (1990) is a wonderful account of the artistic, social, and political interests of the logical positivists and the links between these interests and their philosophical ideas. Passmore 1966 is a good and accessible survey of philosophical movements and trends in the late nineteenth and early twentieth centuries, including absolute idealism.

Hempel, *Aspects of Scientific Explanation* (1965), is the definitive statement of logical empiricism. His *Philosophy of Natural Science* (1966) is the easy version. Carnap's later lectures have been published as *Introduction to the Philosophy of Science* (1995).

An attempt to revive some logical positivist ideas has recently begun; see, for example, Elliott Sober's forthcoming book *Learning from Logical Positivism*.

3

Induction and Confirmation

3.1 The Mother of All Problems

In this chapter we begin looking at a very important and difficult problem, the problem of understanding how observations can *confirm* a scientific theory. What connection between an observation and a theory makes that observation *evidence for* the theory? In some ways, this has been *the* fundamental problem in the last hundred years of philosophy of science. This problem was central to the projects of logical positivism and logical empiricism, and it was a source of constant frustration for them. And although some might be tempted to think so, this problem does not disappear once we give up on logical empiricism. The problem, in some form or other, arises for nearly everyone.

The aim of the logical empiricists was to develop a *logical* theory of evidence and confirmation, a theory treating confirmation as an abstract relation between sentences. It has become fairly clear that their approach to the problem is doomed. The way to analyze testing and evidence in science is to develop a different kind of theory. But it will take a lot of discussion, in this and later chapters, before the differences between approaches that will and will not work in this area can emerge. The present chapter will mostly look at how the problem of confirmation was tackled in the middle of the twentieth century. And that is a tale of woe.

Before looking at twentieth-century work on these issues, we must again look further into the past. The confirmation of theories is closely connected to another classic issue in philosophy: *the problem of induction*. What reason do we have for expecting patterns observed in our past experience to hold also in the future? What justification do we have for using past observations as a basis for generalization about things we have not yet observed?

The most famous discussions of induction were written by the eighteenth-century Scottish empiricist David Hume ([1739] 1978). Hume asked, What reason do we have for thinking that the future will resemble the past? There

is no *contradiction* in supposing that the future could be totally unlike the past. It is *possible* that the world could change radically at any point, rendering previous experience useless. How do we know this will not happen? We might say to Hume that when we have relied on past experience before, this has turned out well for us. But Hume replies that this is begging the question—presupposing what has to be shown. Induction has worked in the past, sure, but that's the *past!* We have successfully used “past pasts” to tell us about “past futures.” But our problem is whether *anything* about the past gives us good information about what will happen *tomorrow*.

Hume concluded that we have no reason to expect the past to resemble the future. Hume was an “inductive skeptic.” He accepted that we all use induction to make our way around the world. And he was not suggesting that we stop doing so (even if we could). Induction is psychologically natural to us. Despite this, Hume thought it had no rational basis. Hume's inductive skepticism has haunted empiricism ever since. The problem of confirmation is not the same as the classical problem of induction, but it is closely related.

3.2 Induction, Deduction, Confirmation, and Explanatory Inference

The logical empiricists tried to show how observational evidence could provide support for a scientific theory. The idea of “support” is important here; there was no attempt to show that scientific theories could be *proved*. Error is always possible, but evidence can support one theory over another.

The cases that were to be covered by this analysis included the simplest and most traditional cases of induction: if we see a multitude of cases of white swans, and no other colors, why does that give us reason to believe that all swans are white? But obviously not all cases of evidence in science are like this. The observational support for Copernicus's theory that the earth goes around the sun, or for Darwin's theory of evolution, seems to work very differently. Darwin did *not* observe a set of individual cases of evolution and then generalize.

The logical empiricists wanted a theory of evidence, or “theory of confirmation,” that would cover all these cases. They were not trying to develop a *recipe* for confirming theories. Rather, the aim was to give an account of the relationships between the statements that make up a scientific theory and statements describing observations, which make the observations support the theory. You might wonder, at this point, what use there could be for a theory with so distant a relationship to actual scientific behavior. Who cares whether a logical analysis of this kind exists or not? In defense of logical empiricism, we might say this: although scientific behavior is not being

directly described by the theory of confirmation, nonetheless scientific procedures might be *based* on assumptions described in the theory of confirmation. Perhaps scientists do many things that cannot be justified if confirmation does not exist.

Let us look more closely at what the logical empiricists tried to do. First, I should say more about the distinction between deductive and inductive logic (a distinction introduced in chapter 2). Deductive logic is the well-understood and less controversial kind of logic. It is a theory of patterns of argument that transmit truth with certainty. These arguments have the feature that *if* the premises of the argument are true, the conclusion is guaranteed to be true. An argument of this kind is *deductively valid*. The most famous example of a logical argument is a deductively valid argument:

PREMISES	All men are mortal. Socrates is a man.
	—————
CONCLUSION	Socrates is mortal.

A deductively valid argument might have false premises. In that case the conclusion might be false as well (although it also might not be). What you get out of a deductive argument depends on what you put in.

The logical empiricists loved deductive logic, but they realized that it could not serve as a complete analysis of evidence and argument in science. Scientific theories do have to be logically *consistent*, but this is not the whole story. Many inferences in science are not deductively valid and give no guarantee. But they still can be *good* inferences; they can still provide *support* for their conclusions.

For the logical empiricists, there is a reason why so much inference in science is not deductive. As empiricists, they believed that all our evidence derives from observation. Observations are always of *particular* objects and occurrences. But the logical empiricists thought that the great aim of science is to discover and establish *generalizations*. Sometimes the aim was seen as describing “laws of nature,” but this concept was also regarded with some suspicion. The key idea was that science aims at formulating and testing generalizations, and these generalizations were seen as having an infinite range of application. No finite number of observations can conclusively establish a generalization of this kind, so these inferences from observations in support of generalizations are always nondeductive. (In contrast, all it takes is *one* case of the right kind to prove a generalization to be *false*; this fact will loom large in the next chapter.)

In many discussions of these topics, the logical empiricists (and some

later writers) used a simple terminology in which all arguments are either deductive or inductive. Inductive logic was thought of as a theory of *all* good arguments that are not deductive. Carnap, especially, used “induction” in a very broad way. But this terminology can be misleading, and I will set things up differently.

I will use the term “induction” only for inferences from particular observations in support of generalizations. To use the most traditional example, the observation of a large number of white swans (and no swans of any other color) might be used to support the hypothesis that all swans are white. We could express the premises with a list of particular cases—“Swan 1 observed at time t_1 was white; swan 2 observed at time t_2 was white. . . .” Or we might simply say: “All the many swans observed so far have been white.” The conclusion will be the claim that all swans are white—a conclusion that could well be false but which is supported, to some extent, by the evidence. Sometimes “enumerative induction” or “simple induction” is used for inductive arguments of this most traditional and familiar kind. Not all inferences from observations to generalizations have this very simple form, though. (And a note to mathematicians: *mathematical induction* is really a kind of *deduction*, even though it has the superficial form of *induction*.)

A form of inference closely related to induction is *projection*. In a projection, we infer from a number of observed cases to arrive at a prediction about the *next* case, not to a generalization about all cases. So we see a number of white swans and infer that the next swan will be white. Obviously there is a close relationship between induction and projection, but (surprisingly, perhaps) there are a variety of ways of understanding this relationship.

Clearly there are other kinds of nondeductive inference in science and everyday life. For example, during the 1980s Luis and Walter Alvarez began claiming that a huge meteor had hit the earth about 65 million years ago, causing a massive explosion and dramatic weather changes that coincided with the extinction of the dinosaurs (Alvarez et al. 1980). The Alvarez team claimed that the meteor caused the extinctions, but let’s leave that aside here. Consider just the hypothesis that a huge meteor hit the earth 65 million years ago. A key piece of evidence for this hypothesis is the presence of unusually high levels of some rare chemical elements, such as iridium, in layers in the earth’s crust that are about 65 million years old. These chemical elements tend to be found in meteors in much higher concentrations than they are near the surface of the earth. This observation is taken to be strong evidence supporting the Alvarez theory that a meteor hit the earth around that time.

If we set this case up as an argument, with premises and a conclusion,

it clearly is not an induction or a projection. We are not inferring to a generalization, but to a hypothesis about a structure or process that would explain the data. A variety of terms are used in philosophy for inferences of this kind. C. S. Peirce called these “abductive” inferences as opposed to inductive ones. Others have called them “explanatory inductions,” “theoretical inductions,” or “theoretical inferences.” More recently, many philosophers have used the term “inference to the best explanation” (Harman 1965; Lipton 1991). I will use a slightly different term—“explanatory inference.”

So I will recognize two main kinds of nondeductive inference, *induction* and *explanatory inference* (plus *projection*, which is closely linked to induction). The problem of analyzing confirmation, or the problem of analyzing evidence, includes all of these.

How are these kinds of inference related to each other? For logical positivism and logical empiricism, induction is the most fundamental kind of nondeductive inference. Reichenbach claimed that all nondeductive inference in science can be reconstructed in a way that depends only on a form of inference that is close to traditional induction. What looks like an explanatory inference can be somehow broken down and reconstructed as a complicated network of inductions and deductions. Carnap did not make this strong claim, but he did seem to view induction as a *model* for all other kinds of nondeductive inference. Understanding induction was in some sense the key to the whole problem. And the majority of the logical empiricist literature on these topics was focused on induction rather than explanatory inference.

So one way to view the situation is to see induction as fundamental. But it is also possible to do the opposite, to claim that explanatory inference is fundamental. Gilbert Harman argued in 1965 that inductions are justified only when they are explanatory inferences in disguise, and others have followed up this idea in various ways.

Explanatory inference seems much more common than induction within actual science. In fact, you might be wondering whether science contains *any* inductions of the simple, traditional kind. That suspicion is reasonable, but it might go too far. Science does contain inferences that look like traditional inductions, at least on the face of them. Here is one example. During the work that led to the discovery of the structure of DNA by James Watson and Francis Crick, a key piece of evidence was provided by “Chargaff’s rules.” These “rules,” described by Erwin Chargaff in 1947, have to do with the relation between the amounts of the four “bases,” C, A, T, and G, that help make up DNA. Chargaff found that in the DNA samples he analyzed, the amounts of C and G were always roughly the same, and the

amounts of T and A were always roughly the same. This fact about DNA became important in the discussions of how DNA molecules are put together. I called it a "fact" just above, but of course Chargaff in 1947 had not observed all the molecules of DNA that exist, and neither have we. In 1947 Chargaff's claim rested on an induction from a small number of cases (in just eight different kinds of organisms). Today we can give an argument for why Chargaff's rules hold that is not just a simple induction; the structure of DNA explains why Chargaff's rules must hold. But it might appear that, back when the rules were originally discovered, the only reason to take the rules to describe all DNA was inductive.

So it might be a good idea to refuse to treat one of these kinds of inference as "more fundamental" than the other. Maybe there is more than one kind of good nondeductive inference (and perhaps there are others besides the ones I have mentioned). Philosophers often find it attractive to think that there is ultimately just one kind of nondeductive inference, because that seems to be a simpler situation. But the argument from simplicity is unconvincing.

Let us return to our discussion of how the problem was handled by the logical empiricists. They used two main approaches. One was to formulate an inductive logic that looked as much as possible like deductive logic, borrowing ideas from deductive logic whenever possible. That was Carl Hempel's approach. The other approach, used by Rudolf Carnap, was to apply the mathematical theory of probability. In the next two sections of this chapter, I will discuss some famous problems for logical empiricist theories of confirmation. The problems are especially easy to discuss in the context of Hempel's approach, which was simpler than Carnap's. A detailed examination of Carnap is beyond the scope of this book. Through his career, Carnap developed very sophisticated models of confirmation using probability theory applied to artificial languages. Problems kept arising. More and more special assumptions were needed to make the results come out right. There was never a knockdown argument against him, but the project came to seem less and less relevant to real science, and it eventually ran out of steam (Howson and Urbach 1993).

Although Carnap's approach to analyzing confirmation did not work out, the idea of using probability theory to understand confirmation remains popular and has been developed in new ways. Certainly this looks like a good approach; it does seem that observing the raised iridium level in the earth's crust made the Alvarez meteor hypothesis *more probable* than before. In chapter 14 I will describe new ways to use probability theory to understand the confirmation of theories.

Before moving on to some famous puzzles, I will discuss a simple proposal that may have occurred to you.

The term *hypothetico-deductivism* is used in several ways by people writing about science. Sometimes it is used to describe a simple view about testing and confirmation. According to this view, hypotheses in science are confirmed when their logical consequences turn out to be true. This idea covers a variety of cases; the confirmation of a white-swan generalization by observing white swans is one case, and another is the confirmation of a hypothesis about an asteroid impact by observations of the true consequences of this hypothesis.

As Clark Glymour has emphasized (1980), an interesting thing about this idea is that it is hopeless when expressed in a simple way, but something like it seems to fit well with many episodes in the history of science. One problem is that a scientific hypothesis will only have consequences of a testable kind when it is combined with other assumptions, as we have seen. But put that problem aside for a moment. The suggestion above is that a theory is confirmed when a true statement about observables can be derived from it. This claim is vulnerable to many objections. For example, any theory T deductively implies $T\text{-or-}S$, where S is any sentence at all. But $T\text{-or-}S$ can be conclusively established by observing the truth of S . Suppose S is observational. Then we can establish $T\text{-or-}S$ by observation, and that confirms T . This is obviously absurd. Similarly, if theory T implies observation E , then the theory $T\&S$ implies E as well. So $T\&S$ is confirmed by E , and S here could be anything at all. (Note the similarity here to a problem discussed at the beginning of section 2.4.) There are many more cases like this.

The situation is strange, and some readers may feel exasperation at this point. People do often regard a scientific hypothesis as supported when its consequences turn out to be true; this is taken to be a routine and reasonable part of science. But when we try to summarize this idea using simple logic, it seems to fall apart. Does the fault lie with the original idea, with our summary of the idea using basic logic, or with basic logic itself? The logical empiricist response was to hang steadfastly onto the logic, and often to hang onto their translations of ideas about science into a logical framework as well. This led them to question or modify some very reasonable-looking ideas about evidence and testing. But it is hard to work out where the fault really lies.

A related feature of logical empiricism is the use of simplified and artificial cases rather than cases from real science. The logical empiricists sought to strip the problem of confirmation down to its bare essentials, and they saw

these essentials in formal logic. But to many, philosophy of science seemed to be turning into an exercise in “logic-chopping” for its own sake. And as we will see in the next sections, even the logic-chopping did not go well.

Despite this, there is a lot to learn from the problems faced by logical empiricism. Confirmation really *is* a puzzling thing. Let us look at some famous puzzles.

3.3 The Ravens Problem

The logical empiricists put much work into analyzing the confirmation of generalizations by observations of their instances. At this point we will switch birds, in accordance with tradition. How is it that repeated observations of black ravens can confirm the generalization that all ravens are black?

First I will deal with a simple suggestion that will not work. Some readers might be thinking that if we observe a large number of black ravens and no nonblack ones, then at least we are cutting down the number of ways in which the hypothesis that all ravens are black might be wrong. As we see each raven, there is one less raven that might fail to fit the theory. So in some sense, the chance that the hypothesis is true should be slowly increasing. But this does not help much. First, the logical empiricists were concerned to deal with the case where generalizations cover an infinite number of instances. In that case, as we see each raven we are not reducing the number of ways in which the hypothesis might fail. Also, note that even if we forget this problem and consider a generalization covering just a finite number of cases, the kind of support that is analyzed here is a very weak one. That is clear from the fact that we get no help with the problem of *projection*. As we see each raven we know there is one less way for the generalization to be false, but this does not tell us anything about what to expect with the *next* raven we see.

So let us look at the problem differently. Hempel suggested that, as a matter of logic, all observations of black ravens confirm the generalization that all ravens are black. More generally, any observation of an *F* that is also *G* supports the generalization “All *F*’s are *G*.” He saw this as a basic fact about the logic of support.

This looks like a reasonable place to start. And here is another obvious-looking point: any evidence that confirms a hypothesis *H* also confirms any hypothesis that is logically equivalent to *H*.

What is logical equivalence? Think of it as what we have when two sentences say the same thing in different terms. More precisely, if *H* is logically

equivalent to *H**, then it is impossible for *H* to be true but *H** false, or vice versa.

But these two innocent-looking claims generate a problem. In basic logic the hypothesis “All ravens are black” is logically equivalent to “All nonblack things are not ravens.” Let us look at this new generalization. “All nonblack things are not ravens” seems to be confirmed by the observation of a white shoe. The shoe is not black, and it’s not a raven, so it fits the hypothesis. But given the logical equivalence of the two hypotheses, anything that confirms one confirms the other. So the observation of a white shoe confirms the hypothesis that all ravens are black! That sounds ridiculous. As Nelson Goodman (1955) put it, we seem to have the chance to do a lot of “indoor ornithology”; we can investigate the color of ravens without ever going outside to look at one.

This simple-looking problem is hard to solve. Debate about it continues. Hempel himself was well aware of this problem—he is the one who originally thought of it. But there has not been a solution proposed that everyone (or even most people) have agreed upon.

One possible reaction is to accept the conclusion. This was Hempel’s response. Observing a white shoe *does* confirm the hypothesis that all ravens are black, though presumably only by a tiny amount. Then we can keep our simple rule that whenever we have an “All *F*’s are *G*” hypothesis, any observation of an *F* that is *G* confirms it and also confirms everything logically equivalent to “All *F*’s are *G*.” Hempel stressed that, logically speaking, an “All *F*’s are *G*” statement is not a statement about *F*’s but a statement about everything in the universe—the statement that if something is an *F* then it is *G*. We should note that according to this reply, the observation of the white shoe also confirms the hypothesis that all ravens are green, that all aardvarks are blue, and so on. Hempel was comfortable with this situation, but most others have not been.

A multitude of other solutions have been proposed. I will discuss just two ideas, which I regard as being on the right track.

Here is the first idea. Perhaps observing a white shoe or a black raven *may or may not* confirm “All ravens are black.” It depends on other factors. Suppose we know, for some reason, that either (1) all ravens are black and ravens are extremely rare, or else (2) most ravens are black, a few are white, and ravens are common. Then a casual observation of a black raven will support (2), a hypothesis that says that not all ravens are black. If all ravens were black, we should not be seeing them at all. Observing a white shoe, similarly, may or may not confirm a given hypothesis, depending on what else we know. This reply was first suggested by I. J. Good (1967).

Good's move is very reasonable. We see here a connection to the issue of holism about testing, discussed in chapter 2. The relevance of an observation to a hypothesis is not a simple matter of the content of the two statements; it depends on other assumptions as well. This is so even in the simple case of a hypothesis like "All F's are G" and an observation like "Object A is both F and G." Good's point also reminds us how artificially simplified the standard logical empiricist examples are. No biologist would seriously wonder whether seeing thousands of black ravens makes it likely that all ravens are black. Our knowledge of genetics and bird coloration leads us to expect some variation, such as cases of albinism, even when we have seen thousands of black ravens and no other colors.

Here is a second suggestion about the ravens, which is consistent with Good's idea but goes further. Whether or not a black raven or a white shoe confirms "All ravens are black" might depend on the *order* in which you learn of the two properties of the object.

Suppose you hypothesize that all ravens are black, and someone comes up to you and says, "I have a raven behind my back; want to see what color it is?" You should say yes, because if the person pulls out a white raven, your theory is refuted. You need to find out what is behind his back. But suppose the person comes up and says, "I have a black object behind my back; want to see whether it's a raven?" Then it does not matter to you what is behind his back. You think that all ravens are black, but you don't have to think that all black things are ravens. In both cases, suppose the object behind his back is a black raven and he does show it to you. In the first situation, your observation of the raven seems relevant to your investigation of raven color, but in the other case it's irrelevant.

So perhaps the "All ravens are black" hypothesis is only confirmed by a black raven when this observation had the *potential to refute* the hypothesis, only when the observation was part of a genuine test.

Now we can see what to do with the white shoe. You believe that all ravens are black, and someone comes up and says, "I have a white object behind my back; want to see what it is?" You should say yes, because if he has a raven behind his back your hypothesis is refuted. He pulls out a shoe, however, so your hypothesis is OK. Then someone comes up and says, "I have a shoe behind my back; want to see what color it is?" In this case you need not care. It seems that in the first of these two cases; you have gained some support for the hypothesis that all ravens are black. In the second case you have not.

So perhaps some white-shoe observations *do* confirm "All ravens are black," and some black-raven observations *don't*. Perhaps there is only

confirmation when the observations arise during a genuine test, a test that has the potential to disconfirm as well as confirm.

Hempel saw the possibility of a view like this. His responses to Good's argument and to the order-of-observation point were similar, in fact. He said he wanted to analyze a relation of confirmation that exists *just between a hypothesis and an observation itself*, regardless of extra information we might have, and regardless of the order in which observations were made. But perhaps Hempel was wrong here; there is *no such relation*. We cannot answer the question of whether an observation of a black raven confirms the generalization unless we know something about the way the observation was made and unless we make assumptions about other matters as well.

Hempel thought that some observations are just "automatically" relevant to hypotheses, regardless of what else is going on. That is true in the case of the deductive refutation of generalizations; no matter how we come to see a nonblack raven, that is bad news for the "All ravens are black" hypothesis. But what is true for deductive *disconfirmation* is not true for confirmation.

Clearly this discussion of order-of-observation does not entirely solve the ravens problem. *Why* does order matter, for example, and what if both properties are observed at once? I will return to this issue in chapter 14, using a more complex framework. Putting it briefly, we can only understand confirmation and evidence by taking into account the *procedures* involved in generating data. Or so I will argue.

I will make one more comment on the ravens problem. This one is a digression, but it does help illustrate what is going on. In psychology there is a famous experiment called the "selection task" (Wason and Johnson-Laird 1972). The experiment has been used to show that many people (including highly educated people) make bad logical errors in certain circumstances. The experimental subject is shown four cards with half of each card masked. The subject is asked to answer this question: "Which masks do you have to remove to know whether it is true that if there is a circle on the left of a card, there is a circle on the right as well?" See fig. 3.1 and try to answer the question yourself before reading the next paragraph.

Large majorities of people in many (though not all) versions of this experiment give the wrong answer. Many people tend to answer "only card A" or "card A and card C." The right answer is A and D. Compare this to the ravens problem; the problems have the same structure. I am sure Hempel would have given the right answer if he had been a subject in the four-card experiment, but the selection task might show something interesting about




Fig. 3.1
The Wason selection task

why confirmation has been hard to analyze. For some reason it is difficult for people to see the importance of “card D” tests in cases like this, and it is easy for people to wrongly think that “card C” tests are important. If you are investigating the hypothesis that all ravens are black, card D is analogous to the situation when someone says he has a white object behind his back. Card C is analogous to the situation where he says he has a black object behind his back. Card D is a real test of the hypothesis, but card C is not. Unmasking Card C is evidentially useless, even though it may fit with what the hypothesis says. Not all observations of cases that fit a hypothesis are useful as tests.

3.4 Goodman’s “New Riddle of Induction”

In this section I will describe an even more famous problem, revealed by Nelson Goodman (1955). This argument looks strange, and it is easy to misinterpret. But the issues it raises are very deep.

First we need to be clear about what Goodman was trying to do with his argument. His primary goal was to show that there cannot be a purely “formal” theory of confirmation. He does not think that confirmation is impossible, or that induction is a myth. He just thinks they work differently from the way many philosophers—especially logical empiricists—have thought.

What is a “formal” theory of confirmation? The easiest way to explain

this is to look at deductive arguments. Recall the most famous deductively valid argument:

Argument 1
 PREMISES All men are mortal.
 Socrates is a man.

 CONCLUSION Socrates is mortal.

The premises, if they are true, guarantee the truth of the conclusion. But the fact that the argument is a good one does not have anything in particular to do with Socrates or manhood. Any argument that has the same form is just as good. That form is as follows:

All F’s are G.
 a is an F.

 a is G.

Any argument with this form is deductively valid, no matter what we substitute for “F,” “G,” and “a.” As long as the terms we substitute pick out definite properties or classes of objects, and as long as the terms retain the same meaning all the way through the argument, the argument will be valid.

So the deductive validity of arguments depends only on the form or pattern of the argument, not the content. This is one of the features of deductive logic that the logical empiricists wanted to build into their theory of induction and confirmation. Goodman aimed to show that this is impossible; there can never be a formal theory of induction and confirmation.

How did Goodman do it? Consider argument 2.

Argument 2
 All the many emeralds observed, in diverse circumstances,
 prior to 2010 A.D. have been green.

All emeralds are green.

This looks like a good inductive argument. (Like some of the logical empiricists, I use a double line between premises and conclusion to indicate that the argument is not supposed to be deductively valid.) The argument does not give us a guarantee; inductions never do. And if you would prefer

to express the conclusion as “probably, all emeralds are green” that will not make any difference to the rest of the discussion.

(If you know something about minerals, you might object that emeralds are regarded as green by definition: emeralds are beryl crystals made green by trace amounts of chromium. Please just regard this as another unfortunate choice of example by the literature.)

Now consider argument 3:

Argument 3

All the many emeralds observed, in diverse circumstances, prior to 2010 A.D. have been grue.

All emeralds are grue.

Argument 3 uses a new word, “grue.” We define “grue” as follows:

GRUE: An object is *grue* if and only if it was first observed before 2010 A.D. and is green, or if it was not first observed before 2010 A.D. and is blue.

The world contains lots of grue things; there is nothing strange about grue objects, even though there is something strange about the word. The grass outside my door as I write this is grue. The sky outside on July 1, 2020, will be grue, if it is a clear day. An individual object does *not* have to change color in order to be grue—this is a common misinterpretation. Anything green that has been observed before 2010 passes the test for being grue. So, all the emeralds we have seen so far have been grue.

Argument 3 does *not* look like a good inductive argument. Argument 3 leads us to believe that emeralds observed in the future will be blue, on the basis of previously observed emeralds being green. The argument also conflicts with argument 2, which looks like a good argument. But arguments 2 and 3 have *exactly the same form*. That form is as follows:

All the many E's observed, in diverse circumstances, prior to 2010 A.D., have been G.

All E's are G.

We could represent the form even more schematically than this, but that does not matter to the point. Goodman's point is that two inductive arguments can have the exact same form, but one argument can be good while

the other is bad. So what makes an inductive argument a good or bad one cannot be just its form. Consequently, there can be no purely formal theory of induction and confirmation. Note that the word “grue” works perfectly well in *deductive* arguments. You can use it in the form of argument 1, and it will cause no problems. But induction is different.

Suppose Goodman is right, and we abandon the idea of a formal theory of induction. This does not end the issue. We still need to work out *what* exactly is wrong with argument 3. This is the new riddle of induction.

The obvious thing to say is that there is something wrong with the word “grue” that makes it inappropriate for use in inductions. So a good theory of induction should include a *restriction* on the terms that occur in inductive arguments. “Green” is OK and “grue” is not.

This has been the most common response to the problem. But as Goodman says, it is very hard to spell out the details of such a restriction. Suppose we say that the problem with “grue” is that its definition includes a reference to a specific time. Goodman's reply is that whether or not a term is defined in this way depends on which language we take as our starting point. To see this, let us define a new term, “bleen.”

BLEEN: An object is *bleen* if and only if it was first observed before 2010 A.D. and is blue, or if it was not first observed before 2010 A.D. and is green.

We can use the English words “green” and “blue” to define “grue” and “bleen,” and if we do so we must build a reference to time into the definitions. But suppose we spoke a language that was like English except that “grue” and “bleen” were basic, familiar terms and “green” and “blue” were not. Then if we wanted to define “green” and “blue,” we would need a reference to time.

GREEN: An object is *green* if and only if it was first observed before 2010 A.D. and is grue, or if it was not first observed before 2010 A.D. and is bleen.

(You can see how it will work for “blue.”) So Goodman claimed that whether or not a term “contains a reference to time” or “is defined in terms of time” is a *language-relative* matter. Terms that look OK from the standpoint of one language will look odd from another. So if we want to rule out “grue” from inductions because of its reference to time, then whether an induction is good or bad will depend on *what language we treat as our starting point*. Goodman thought this conclusion was fine. A good induction, for Goodman, must use terms that have a history of normal use in our

community. That was his own solution to his problem. Most other philosophers did not like this at all. It seemed to say that the value of inductive arguments depended on irrelevant facts about which language we happen to use.

Consequently, many philosophers have tried to focus not on the words “green” and “grue” but on the *properties* that these words pick out, or the *classes* or *kinds* of objects that are grouped by these words. We might argue that greenness is a natural and objective feature of the world, and grueness is not. Putting it another way, the green objects make up a “natural kind,” a kind unified by real similarity, while the grue objects are an artificial or arbitrary collection. Then we might say: a good induction has to use terms that we have reason to believe pick out natural kinds. Taking this approach plunges us into hard problems in other parts of philosophy. What *is* a property? What *is* a “natural kind”? These are problems that have been controversial since the time of Plato.

Although Goodman’s problem is abstract, it has interesting links to real problems in science. In fact, Goodman’s problem encapsulates within it several distinct hard methodological issues in science; that is partly why the problem is so interesting. First, there is a connection between Goodman’s problem and the “curve-fitting problem” in data analysis. Suppose you have a set of data points in the form of x and y values, and you want to discern a general relationship expressed by the points by fitting a function to them. The points in figure 3.2 fall almost exactly on a straight line, and that seems to give us a natural prediction for the y value we expect for $x = 4$. However, there is an infinite number of different mathematical functions that fit our three data points (as well or better) but which make different predictions for the case of $x = 4$. How do we know which function to use? Fitting a strange function to the points seems to be like preferring a grue induction over a green induction when inferring from the emeralds we have seen.

Scientists dealing with a curve-fitting problem like this may have extra information telling them what sort of function is likely here, or they may prefer a straight line on the basis of *simplicity*. That suggests a way in which we might deal with Goodman’s original problem. Perhaps the green induction is to be preferred on the basis of its simplicity?

That might work, but there are problems. First, is it really so clear that the green induction is simpler? Goodman will argue that the simplicity of an inductive argument depends on which language we assume as our starting point, for the kinds of reasons given earlier in this section. For Goodman, what counts as a simple pattern depends on which language you speak or which categorization you assume. Also, though a preference for




Fig. 3.2
The curve-fitting problem

simplicity is very common in science, such a preference is often hard to justify. Simpler theories are easier for us to work with, but that does not seem to give us reason to prefer them if we are seeking to learn what the world is really like. Why should the world be simple rather than complex?

Earlier I mentioned attempts to solve Goodman’s problem using the idea of a “natural kind,” a collection unified by real similarity as opposed to stipulation or convention. Though this term is philosophical, a lot of argument within science is concerned just this sort of problem—with getting the right *categories* for prediction and extrapolation. The problem is especially acute in sciences like economics and psychology that deal with complex networks of similarities and differences across the cases they try to generalize about. Do all economies with very high inflation fall into a natural kind that can be used to make general predictions? Are the mental disorders categorized in psychiatric reference books like the *DSM IV* really natural kinds, or have we applied standard labels like “schizophrenia” to groups of cases that have no real underlying similarity? The periodic table of elements in chemistry seems to pick out a set of real natural kinds, but is this something we can hope for in all sciences? If so, what does that tell us about inductive arguments in different fields?

That concludes our initial foray into the problems of induction and confirmation. These problems are simple, but they are very resistant to solution. For a good part of the twentieth century, it seemed that even the most innocent-looking principles about induction and confirmation led straight into trouble.

Later (especially in chapter 14) I will return to these problems. But in the next chapter we will look at a philosophy that gets a good part of its motivation from the frustrations discussed in this chapter.

Further Reading

Once again, Hempel's *Aspects of Scientific Explanation* (1965) is a key source, containing a long (and exhausting) chapter on confirmation. Skyrms, *Choice and Chance* (2000), is a classic introductory book on these issues, and it introduces probability theory as well. Even though it argues for a view that will not be discussed until chapter 14, Howson and Urbach's *Scientific Reasoning* (1993) is a useful introduction to various approaches to confirmation. It has the most helpful short summary of Carnap's ideas that I have read. Carnap's magnum opus on these issues is his *Logical Foundations of Probability* (1950). For a discussion of explanatory inference, see Lipton, *Inference to the Best Explanation* (1991).

For the use of order-of-observation to address the ravens problem, see Horwich, *Probability and Evidence* (1982), but you should probably read chapter 14 of this book first.

Goodman's most famous presentation of his "new riddle of induction" is in *Fact, Fiction & Forecast* (1955). The problem is in chapter 3 (along with other interesting ideas), and his solution is in chapter 4. His subsequent papers on the topic are collected in *Problems and Projects* (1972). Douglas Stalker has edited a collection on Goodman's riddle, called *Grue!* (1994). It includes a very detailed bibliography. The Quine and Jackson papers are particularly good.

For discussions of properties and kinds, and their relevance to induction, see Armstrong 1989, Lewis 1983, Dupre 1993, and Kornblith 1993. (These are fairly advanced discussions, except for Armstrong's, which is introductory.) There is a good discussion of simplicity in Sober 1988.

4

Popper: Conjecture and Refutation

4.1 Popper's Unique Place in the Philosophy of Science

Karl Popper is the only philosopher discussed in this book who is regarded as a hero by many scientists. Attitudes toward philosophy among scientists vary, but hardly ever does a philosopher succeed in *inspiring* scientists in the way Popper has. It is also rare for a philosopher's view of science to be used within a scientific debate to justify one position over another. This has happened with Popper too. Within biology, recent debates about the classification of organisms and about ecology have both seen Popper's ideas used in this way (Hull 1999). I once went to a lecture by a famous virologist who had won a Nobel Prize in medicine, to hear about his work. What I heard was mostly a lecture about Popper. In 1965, Karl Popper even became *Sir* Karl Popper, knighted by the queen of England.

Popper's appeal is not surprising. His view of science is centered around a couple of simple, clear, and striking ideas. His vision of the scientific enterprise is a noble and heroic one. Popper's theory of science has been criticized a great deal by philosophers over the years. I agree with many of these criticisms and don't see any way for Popper to escape their force. Despite the criticism, Popper's views continue to have an important place in philosophy and continue to appeal to many working scientists.

4.2 Popper's Theory of Science

Popper began his intellectual career in Vienna, between the two world wars. He was not part of the Vienna Circle, but he did have contact with the logical positivists. This contact included a lot of disagreement, as Popper developed his own distinctive position. Popper does count as an "empiricist" in the broad sense used in this book, but he spent a lot of time distinguishing his views from more familiar versions of empiricism. Like the logical positivists, Popper left Europe upon the rise of Nazism, and after