Chapter 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 fairly rare for a philosopher's view of science to be used within a scientific debate to justify one position over another. This too has happened with Popper. Within biology, debates about the classification of organisms and about ecology have seen Popper's ideas used in this way (Hull 1999). In 1965, Karl Popper became Sir Karl, 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 rather 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—he saw some things that others seemed not to see—and his ideas 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. Like the logical positivists, Popper left Europe upon the rise of Nazism, and after spending the war years in New Zealand, he moved to the London School of Economics, where he remained for the rest of his career. There he built a loyal group of allies, whom he often accused of disloyalty. His seminar series at the London School of Economics became famous for its grueling questioning, and for the fact that speakers had a difficult time presenting much of their lectures because of Popper's interruptions. Popper once had a famous confrontation with Wittgenstein, on the latter's turf at Cambridge University. One version of the story, told by Popper himself, has Wittgenstein brandishing a fireplace poker during a discussion of ethical rules, leading Popper to give an example of an ethical rule: "not to threaten visiting lecturers with pokers." Wittgenstein stormed out. Other versions of the story, including those told by Wittgenstein's allies, deny Popper's account.

The logical positivists developed their theory of science as part of a general theory of language, meaning, and knowledge. Popper was not much interested in these broader topics, at least initially; his primary aim was to understand science. As his first order of business, he wanted to understand the difference between scientific and nonscientific theories. In particular, he wanted to distinguish science from "pseudo-science." He did not regard pseudo-scientific ideas as meaningless; they just weren't science. For Popper, an inspiring example of genuine science was the work of Einstein. Examples of pseudo-science were Freudian psychology and Marxist views about society and history.

Popper called the problem of distinguishing science from non-science the "problem of demarcation." All of Popper's philosophy starts from his proposed solution to this problem. "Falsificationism" was the name Popper gave to his solution. Falsificationism claims that *a hypothesis is scientific if and only if it has the potential to be refuted by some possible observation*. To be scientific, a hypothesis has to take a risk, has to "stick its neck out." If a theory takes no risks at all, because it is compatible with every possible observation, then it is not scientific. As I said above, Popper held that Marx's and Freud's theories were not scientific in this sense. No matter what happens, Popper thought, a Marxist or Freudian can fit it into their theory somehow. These theories are never exposed to any risks.

So far I have described Popper's use of falsifiability to distinguish scientific from nonscientific theories. Popper also made use of the idea of falsification in a more far-reaching way. He claimed that all testing in science has the form of attempting to refute theories by means of observation. Crucially, for Popper it is never possible to confirm or establish a theory by showing its agreement with observations. Confirmation is a myth. The only thing an observational test can do is to show that a theory is false. So the truth of a scientific theory can never be supported by observational evidence, not even a little bit, and not even if the theory makes a huge number of predictions that all come out as expected.

As you might expect, Popper was a severe critic of the logical empiricists' attempts to develop a theory of confirmation or inductive logic. The problems they encountered, some of which I discussed in chapter 3, were music to his ears. Popper, like Hume, was an inductive skeptic, and Popper was skeptical (at least officially) about all forms of confirmation and support other than deductive logic itself.

Skepticism about induction and confirmation is a much more controversial position than Popper's use of falsification to solve the demarcation problem. Most philosophers of science have thought that if induction and confirmation are myths, that is very bad news for science. Popper tried to argue that there is no reason to worry; induction is a myth, but science does not need it anyway. Inductive skepticism, for Popper, is no threat to the rationality of science. In the opinion of most philosophers, Popper's attempt to defend this radical claim was not successful, and some of his discussions of this topic are rather misleading to readers. Some of the scientists who regard Popper as a hero do not realize that Popper believed it is never possible to confirm a theory, not even slightly, and no matter how many observations the theory helps us to predict successfully.

Popper placed great emphasis on the idea that we can never be completely sure that a theory is true. After all, Newton's physics was viewed as the best-supported theory in science, but early in the twentieth century it was shown to be false in several respects. However, almost all philosophers of science accept that we can never be 100 percent certain about factual matters, especially those discussed in science. This position, that we can never be completely certain about factual issues, is often known as *fallibilism* (a term due to C. S. Peirce). Most philosophers of science accept fallibilism. The harder question is whether or not we can be reasonable in increasing our confidence in the truth of a theory when it passes observational tests. Popper said no. The logical empiricists and most other philosophers of science say yes.

Popper had a fairly simple view of how testing in science proceeds. We take a theory that someone has proposed, and deduce an observational prediction from it. We then check to see if the prediction comes out as the theory says it will. If the prediction fails, then we have refuted, or falsified, the theory. If the prediction does come out as predicted, then all we should say is that we have not yet falsified the theory. For Popper, we cannot conclude that the theory is true, or that it is probably true, or even that it is more likely to be true than it was before the test. The theory *might* be true, but we can't say more than that.

We then try to falsify the theory in some other way, with a new prediction. We keep doing this until we have succeeded in falsifying it. What if years pass and we seem to never be able to falsify a theory, despite repeated tests? We can say that the theory has now survived repeated attempts to falsify it, but that's all. We never increase our confidence in the truth of the theory, and ideally, we never stop trying to falsify it. That's not to say we should spend all our time testing theories that have passed tests over and over again. We do not have the time and resources to test everything that could be tested. But that is just a practical constraint. According to Popper, we should always retain a tentative attitude toward our theories, no matter how successful they have been in the past.

In defending this view, Popper made much of the difference between confirming and disconfirming statements of scientific law. If someone proposes a law of the form "All *F*s are *G*," all it takes is one observation of an *F* that is not a *G* to falsify the hypothesis. This is a matter of deductive logic. But it is never possible to assemble enough observations to conclusively demonstrate the truth of such a hypothesis. You might wonder about situations where there are only a small number of *F*s and we could hope to check them all. Popper and the logical empiricists regarded these as unimportant situations that do not often arise in science. Their aim was to describe testing in situations where there is a huge or infinite number of cases covered by a hypothesized law or generalization. So Popper stressed that universal statements are hard or impossible to verify but easy, in principle, to falsify. The logical empiricist might reply that statements of the form "Some *F*s are *G*" have the opposite feature; they are easy to verify but hard or impossible to falsify. But

Popper claimed (and the logical empiricists tended to agree) that real scientific theories rarely take this form, even though some statements in science do.

Despite insisting that we can never support or confirm scientific theories, Popper believed that science is a search for true descriptions of the world. How can one search for truth if confirmation is impossible? This is an unusual kind of search. We might compare it to a certain kind of search for the Holy Grail, conducted by an imaginary medieval knight.

Suppose there are lots of grails around, but only one of them is holy. In fact, the number of non-holy grails is infinite or enormous, and you will never encounter them all in a lifetime. All the grails glow, but only the Holy Grail glows forever. The others eventually stop glowing, but there is no telling when any particular non-holy grail will stop glowing. All you can do is pick up one grail and carry it around and see if it keeps on glowing. You are only able to carry one at a time. If the one you are carrying is the Holy Grail, it will never stop glowing. But you would never know if you currently had the Holy Grail, because the grail you are carrying might stop glowing at any moment. All you can do is reject grails that are clearly not holy (since they stop glowing at some point) and keep picking up a new one. You will eventually die (with no afterlife, in this scenario) without knowing whether you succeeded.

This is similar to Popper's picture of science's search for truth. All we can do is try out one theory after another. A theory that we have failed to falsify up till now might be true. But if so, we will never know this or even have reason to increase our confidence.

Here is a quote from Popper's most important work, *The Logic of Scientific Discovery*, expressing his view:

> I think that we shall have to get accustomed to the idea that we must not look upon science as a "body of knowledge," but rather as a system of hypotheses; that is to say, as a system of guesses or anticipations which in principle cannot be justified, but with which we work as long as they stand up to tests, and of which we are never justified in saying that we know they are "true" or "more or less certain" or even "probable." (1959, 317)

4.3 *Popper on Scientific Change*

So far I have described Popper's views about the demarcation of science from non-science and the nature of scientific testing. Popper also used the idea of falsification to propose a theory of scientific change.

Popper's theory has an appealing simplicity. Science changes via a two-step cycle that repeats endlessly. Stage 1 in the cycle is *conjecture*—a scientist will offer a hypothesis that might describe and explain some part of the world. A good conjecture is a bold one, one that takes a lot of risks by making novel predictions. Stage 2 in the cycle is *attempted refutation*—the hypothesis is subjected to critical testing, in an attempt to show that it is false. Once the hypothesis is refuted, we go back to stage 1 again—a new conjecture is offered. That is followed by stage 2, and so on.

As the process moves along, it is natural for a scientist to propose conjectures that have some relation to previous ones. A theoretical idea can be refined and modified via many rounds of conjecture and refutation. That is fine, for Popper, though it is not essential. One thing that a scientist should not do, however, is react to the falsification of one conjecture by cooking up a new conjecture that is designed to just avoid the problems revealed by earlier testing, and that goes no further. We should not make ad hoc moves that merely patch the problems found in earlier conjectures. Instead, a scientist should constantly strive to increase the breadth of application of a theory, and increase the precision of its predictions. That means constantly trying to increase the "boldness" of conjectures.

What sort of theory is this? Popper intended it as a description of the general pattern that we actually see in science, and as a description of good scientific behavior as well. He accepted that not all scientists succeed in sticking to this pattern of behavior all the time. Sometimes people become too wedded to their hypotheses and refuse to give them up when testing tells them to. But Popper thought that a lot of actual scientific behavior does follow this pattern and that we see it especially in great scientists such as Einstein. For Popper, a good or great scientist is someone who combines two features, one corresponding to each stage of the cycle. The first feature is an ability to come up with imaginative, creative, and risky ideas. The second is a hard-headed willingness to subject these imaginative ideas to rigorous critical testing. A good scientist has a creative, almost artistic, streak and a tough-minded, no-nonsense streak. Imagine a hard-headed cowboy out on the range, with a Stradivarius violin in his saddlebags. (Perhaps at this point you can see some of the reasons for Popper's popularity among scientists.)

Popper's view on this issue can apparently be applied in the same way to individuals and to groups of scientists. An isolated individual can behave scientifically by engaging in the process of conjecture and refutation. And a collection of scientists can each, at an individual level, follow Popper's two-step procedure. But it seems that another possibility is a division of labor; one individual (or team) comes up with a conjecture, and another attempts to refute it. The basic idea of a conjecture-and-refutation pattern seems compatible with all these possibilities. But the case where individual *A* does the conjecture and individual *B* does the refutation will be suspicious to Popper. If individual *A* is a true scientist, they should take a critical attitude toward their own ideas. If individual *A* is completely fixated on their conjecture, and individual *B* is fixated on showing that *A* is wrong in order to advance a different conjecture, this is not good scientific behavior, according to Popper.

This raises an interesting question. Empiricist philosophies emphasize the virtues of open-mindedness, and Popper's view, which I see as an unorthodox version of empiricism, is no exception. But perhaps an open-minded community can be built out of a collection of rather closed-minded individuals. If actual scientists are wedded to their own conjectures, but each is wedded to a different conjecture and would like to prove the others wrong, shouldn't the overall process of conjecture and refutation work? What is wrong with the situation where *B*'s role is to critically test *A*'s ideas? So long as the testing occurs, what does it matter whether *A* or *B* does it? One possible problem here is that if everyone is closed-minded in this way, the results of tests might have no impact on what people believe. Real openness in the community would require that falsifications retain their bite once they have been achieved.

We can also imagine another kind of division of labor, with specialists on each side of the Popperian combination—specialist conjecturers and refuters. To some extent, we do surely see something like this. All this presses the following question on Popper: if a combination of conjecture and refutation is what is characteristic of science, then why isn't it enough to have some combination of socially organized behaviors that gives rise to openness in the community as a whole?

Although Popper did take an interest in community standards within science, he did seem to have a picture in which the good scientist should, as an individual, be willing to perform both the imaginative and the critical roles. Good scientists should retain a tentative attitude toward all theories, including their own: "whenever we propose a solution to a problem, we ought to try as hard as we can to overthrow our solution, rather than defend it" (1959, p. xix).

This relates to the inspirational role of Popper's ideas in science, which continues to this day. A few years ago I was driving down a highway and heard on the radio part of a long interview with a successful biologist looking back on his career, talking about how he came to do work that made a difference. He gave a lot of credit to Popper for getting him to think the right way about science. Here is what he emphasized most. Popper taught him that it is *OK to be wrong*. It is good to be wrong, if you and others can learn from the error. He came back to this idea several times in the interview; it really seemed to affect how he approached his work.

I agree that Popper was onto something here, and there seems to be something healthy in the Popperian attitude. "Seems," I said—so far this is just an impression. Soon in this book we'll encounter arguments that this attitude might not be as helpful for science as it appears.

I will make one more point before moving on to criticisms of Popper. The two-step process of conjecture and refutation that Popper describes has a striking resemblance to another two-step process: Darwin's explanation of biological evolution in terms of *variation* and *natural selection*. In science, according to Popper, scientists generate conjectures that are subjected to critical testing. In evolution, according to both Darwin himself and more recent versions of evolutionary theory, populations evolve via a process in which variations appear in organisms in a random or "undirected" way, and these novel characteristics are tested through their effects on the organism in its interactions with the environment. Variations that help organisms to survive and reproduce, and that are of the kind that gets passed on in reproduction, tend to be preserved and become more common in the population over time.

Ironically, at one time Popper thought that Darwinism was not a scientific theory, but he later retracted that claim. In any case, Popper and others have explored in detail the analogy between Popperian science and Darwinian evolution. The analogy should not be taken too seriously; evolution is not a process in which populations really "search" for anything, in the way that scientists search for good theories, and there are other important differences too. But the similarity is certainly interesting.

4.4 *Objections to Popper on Falsification*

I now turn to a critical assessment of Popper's ideas, beginning with his solution to the demarcation problem. Is falsifiability a good way to distinguish scientific ideas from nonscientific ones?

Let me first say that I think this question probably has no answer in the form in which Popper expressed it. We should not expect to be able to go through a list of statements or theories and label them "scientific" or "not scientific." However, I suggest that something fairly similar to Popper's question about demarcation does make sense: can we describe a distinctive scientific *strategy* of investigating the world, a scientific way of *handling* ideas?

Some of Popper's ideas are useful in trying to answer this question. In particular, Popper's claim that scientific theories should take *risks* is a good one. But Popper had an overly simple picture of how this risk-taking works. For Popper, theories have the form of generalizations, and they take risks by prohibiting certain kinds of particular events from being observed. If we believe that all pieces of iron, of whatever size and shape, expand when heated, then our theory forbids the observation of something that we know to be a piece of iron contracting when heated. A problem may have occurred to you: how sure can we be that, if we see a piece of "iron" contracting when heated, that it is really iron? We might also have doubts about our measurements of the contraction and the temperature change. Maybe the generalization about iron expanding when heated is true, but our assumptions about the testing situation and our ability to know that a sample is made of iron are false.

This problem is a reappearance of an issue discussed in chapter 2: holism about testing. Whenever we try to test a theory by comparing it with observations, we must make a large number of additional assumptions in order to bring the theory and the observations into contact with each other. If we want to test whether iron always expands when heated, we need to make assumptions about our ability to find or make reasonably pure samples of iron. If we want to test whether the amounts of the bases C and G are equal and the amounts of A and T are equal in all samples of DNA (Chargaff's rules, from chapter 3), we need to make a lot of assumptions about our chemical techniques. If we observe an unexpected result (iron contracting on heating, twice as much C as G in a sample of DNA), it is always possible to blame one of these extra assumptions rather than the theory we are trying to test. In extreme cases, we might even claim that the apparent observation was completely misunderstood or wrongly described by the observers. Indeed, this is not so uncommon in our attempts to work out what to make of reports of miracles and UFO abductions. So how can we really use observations to falsify theories in the way Popper wants? This is a problem not just for Popper's solution to the demarcation problem, but for his whole theory of science as well.

Some familiar ways of talking about testing—seen in Popper and also others—can be quite misleading on this front. People often say: a theory implies an observational prediction. Or, as in Popper's demarcation criterion, a hypothesis is only scientific if it has the potential to be refuted by some possible observation that clashes with it. It is never as simple as this, even in the ultra-simple "All ravens are black" cases. (I discussed this also in the optional section at the end of the previous chapter.) A generalization can tell you that if some object is a raven, then it has to be black. But working out whether there are any ravens around is a separate matter. The generalization plus some other assumptions tells you what you should see. And when the observation that comes back is a surprising one, there is always more than one possible thing to blame. A theory can't "take risks," in the way Popper likes, all on its own.

Popper was aware of this problem, and he struggled with it. He regarded the extra assumptions needed to connect theories with testing situations as scientific claims that might well be false—these are conjectures too. We can try to test these conjectures separately. But Popper conceded that logic itself can never force a scientist to give up a particular theory in the face of surprising observations. Logically, it is always possible to blame other assumptions involved in the test. Popper thought that a good scientist would not try to do this; a good scientist would want to expose the theory itself to tests and not try to deflect the blame.

Does this answer the holist objection? What Popper has done is move from describing a characteristic of scientific *theories* to describing a characteristic of scientific *behavior*. In some ways this is a retraction of his initial aim, which was to describe something about scientific theories themselves that makes them special. That is a problem. Then again, this shift to describing scientific modes of thought and behavior, rather than theories, might be a step forward. I'll have a closer look at this idea later in this chapter.

Popper also accepted that we cannot be completely certain about the observation reports that we use to falsify theories. We have to regard the acceptance of an observation report as a "decision," one that is freely made. Once we have made the decision, we can use the observation report to falsify any theory that conflicts with it. But for Popper, any falsification process is based, in the end, on a decision that could be challenged. Someone might come along later and try to show, via more testing, that the observation report was not a good one; this person might investigate whether the conditions of observation were misleading. That testing will have the same conjecture-and-refutation form described earlier. So the investigation into the controversial observation ultimately depends on "decisions" too.

Is this bad news for Popper? Popper insisted that making these decisions about single observations is very different from making free decisions directly about the theories themselves. But what sort of difference is this? If observation reports rest on nothing more than decisions, and these determine our choice of theories, how is that better than directly choosing the theories themselves, without worrying about observation? And couldn't we just "decide" to hang onto a theory and reject the observation reports that conflict with it? I am not saying that we should do these things, just that Popper has not given us a good reason not to do them. I believe that we should not do these things because we have good reason to believe that observation is a generally reliable way of forming beliefs, at least of some particular kinds. As I will argue in chapter 9, we can make use of a scientific theory of perception at this point in the story. But that argument will have to come later. Popper himself does not try to answer these questions by giving an argument about the reliability of perception.

This point about the role of decisions affects Popper's ideas about demarcation as well as his ideas about testing. Any system of hypotheses can be held onto despite apparent falsification if people are willing to make certain decisions. Given this, does Popper's view end up giving us any way to differentiate between science and pseudo-science? The answer is "yes and no." The "no" comes from the fact that scientific theories can be handled in a way that makes them immune to falsification, and nonscientific theories can be rejected if people decide to accept claims about particular matters that are incompatible with the theory. But there is a "yes" part in the answer as well. A scientific theory is falsifiable via a certain *kind* of decision—a decision about an observation report, which together with background assumptions can clash with a theory. A pseudo-scientific theory, Popper says, does not have this feature. So if a pseudo-scientific theory is to be rejected, some different kind of decision must be made.

There is another problem with Popper's views about falsification to discuss. The problem is bad for Popper, but I should emphasize that it is bad for many others as well. What can Popper say about theories that do not claim that some observation is forbidden, but only that it is very *unlikely*? If I believe that a certain coin is "fair," I can deduce from this hypothesis various claims about the probabilities of long all-heads or all-tails sequences of tosses. Suppose I observe a hundred tosses turning up heads a hundred times. This is very unlikely, according to my hypothesis about the coin, but it is not impossible. Any finite stretch of heads tosses is possible with a fair coin, although longer and longer runs of heads are treated by the theory as more and more unlikely. But if a hypothesis does not forbid any particular observations, then, according to Popper, it is taking no risks. This seems to entail that for Popper, theories that ascribe very low probabilities to specific observations, but do not rule them out altogether, are unfalsifiable and hence unscientific.

Popper's response was to accept that, logically speaking, hypotheses of this kind are unscientific. But this seems to make a mockery of the important role of probability in science. So Popper said that a scientist can decide that if a theory claims that a particular observation is extremely improbable, the theory *in practice* rules out that observation. If the observation is made, the theory is, in practice, falsified. According to Popper, it is up to scientists to work out, for their own fields, what sort of probability is so low that events of that kind are treated as prohibited. Probabilistic theories can only be construed as falsifiable in a special "in practice" sense. And we have here another role for decisions in Popper's philosophy of science, as opposed to the constraints of logic.

Popper is right that scientists reject theories when observations occur that the theory says are highly improbable (although which kinds of improbability have this importance is a complicated matter). And Popper is right that scientists spend a good deal of time working out "how improbable is *too* improbable." Complex statistical methods are used to help scientists with these decisions. But in making this move, Popper has damaged his original picture of science. This was a picture in which observations, once accepted, have the power to decisively refute theoretical hypotheses. That is a matter of deductive logic, as Popper endlessly stressed (though, as we saw, this has to work with the aid of background assumptions). Now Popper is saying that falsification can occur without its being backed up by a deductive logical relation between observation and theory.

4.5 *Objections to Popper on Confirmation*

As described earlier, Popper believed that theories can never be confirmed by observations, and he thought that inductive arguments are never justified. Popper thought that a theory of the rational choice of theories could be given entirely in terms of falsification, so he thought that rejecting induction and confirmation was no problem.

In the previous section I discussed problems with Popper's views about falsification. But let us leave those problems aside now, and assume in this section that we can use Popperian falsification as a method for decisively rejecting theories. If we make this assumption, is Popper's attempt to describe rational theory choice successful?

Here is simple problem that Popper has a very difficult time with. Suppose we are trying to build a bridge, and we need to use physical theories to tell us which designs are stable and will support the weight that the bridge must carry. This is a situation in which we must apply our scientific theories to a practical task. As a matter of fact, engineers and scientists in this situation will undoubtedly tend to use physical theories that have survived empirical testing; they will use "tried and true" methods as far as possible. The empiricist approach to the philosophy of science holds that such a policy is rational. A problem for empiricism is to explain in more detail why this policy is the right one. That task is hard, as I hope became clear in chapter 3. But let us focus on Popper, who wants to avoid the need for a theory of confirmation. How does Popper's philosophy treat the bridge-building situation?

Popper can say something about why we should prefer to use a theory that has not been falsified over a theory that has been falsified. Theories that have been falsified have been shown to be false (here again I ignore the problems discussed in the previous section). But suppose we have to choose between (1) a theory that has been tested many times and has passed every test and (2) a brand-new theory that has just been conjectured and has never been tested. Neither theory has been falsified. We would ordinarily think that the rational thing to do is to choose the theory that has survived testing. But what can Popper say about this choice? Why exactly would it be irrational, for Popper, to build the bridge using a new theory that has not yet been tested?

Popper recognized and struggled with this problem too. Perhaps this has been the most common objection to Popper from other empiricist philosophers (e.g., Salmon 1981). Popper is unable to give a very good reply. Popper refuses to say that when a theory passes tests, we have more reason to believe that the theory is true. Both the untested theory and the well-tested theory are just conjectures. But Popper did devise a special concept to use in this situation. He said that a theory that has survived many attempts to falsify it is "corroborated." And when we face choices like the bridge-building one, it is rational to choose corroborated theories over theories that are not corroborated.

What is corroboration? Popper gave a technical definition and held that we can measure the amount of corroboration that a theory has at a particular time. The technicalities do not matter, though. We need to ask, what *sort* of property is corroboration? Has Popper just given a new name to confirmation? If so, he can answer the question about building the bridge, but he has given up one of the main features that distinguishes his view from the logical empiricists and everyone else. If corroboration is totally different from confirmation—so different that we cannot regard corroboration as any guide to a theory's truth—then why should we choose a corroborated theory when we build the bridge?

This issue has been much discussed (see Newton-Smith 1981). Popper's concept of corroboration can be interpreted in a way that makes it different from confirmation, but then Popper can give no good answer to the question of why we should choose corroborated theories over new ones when building bridges.

To understand corroboration, think of the difference between an academic transcript and a letter of recommendation. This distinction should be vivid to students! An academic transcript says what you have done. It measures your past performance, but it does not contain explicit predictions about what you will do in the future. A letter of recommendation usually says something about what you have done, and it also makes claims about how you are likely to do in the future. Confirmation, as understood by the logical empiricists, is something like a letter of recommendation for a scientific theory. Corroboration, for Popper, is like an academic transcript. And Popper thought that no good reasons could be given for believing that past performance is a reliable guide to the future. So corroboration is entirely backward-looking. Consequently, no reason can be given for building a bridge with a corroborated theory rather than a noncorroborated but unfalsified one.

I think the best thing for Popper to say about the bridge-building situation is to stick to his inductive skepticism. He should argue that we really don't know what will happen if we build another bridge with a design that has worked in the past. Maybe it will stay up and maybe it won't. There might also be practical reasons for choosing that design if we are very familiar with it. But if someone comes along with a brand-new, untested design, we won't know whether it's a bad design until we try it. Popper liked to say that there is no alternative policy that is *more* rational than using the familiar and well-tested design, and we do have to make some decision. So we can go ahead and use the established design. But as Wesley Salmon (1981) replied, this does not help at all. If confirmation does not exist, then it seems there is also no policy that is more rational than choosing the *un*tested design. All we have here is a kind of "tie" between the options.

For most people, this is an unsatisfactory place for a philosophy of science to end up. Inductive skepticism of this kind is hard to take seriously outside of abstract, academic discussion. However, the efforts seen over several centuries have shown how hard it is to produce a good theory of induction and confirmation. One of the valuable roles of Popper's philosophy is to show what sort of theory of science might be possible if we give up on induction and confirmation.

In the first chapter of this book, I said that few philosophers still try to give descriptions of a definite "scientific method," where this is construed as something like a recipe for science. Popper is a partial exception here, since he does come close to giving a kind of recipe (although Popper insists there is no recipe for coming up with interesting conjectures). His view has an interesting relationship to descriptions of scientific method given in science textbooks.

In many textbooks, one finds something called the "hypothetico-deductive method." Back in chapter 3, I discussed a view about confirmation that is often called "hypothetico-deductivism." Now we are dealing with a method rather than a theory of confirmation. Science textbooks are more cautious about laying out recipes for science than they used to be, but descriptions of the hypothetico-deductive method are still fairly common. Formulations of the method vary, but some are basically a combination of Popper's view of testing and a less skeptical view about confirmation. In these versions, the hypothetico-deductive method is a process in which scientists come up with conjectures and then deduce observational predictions from those conjectures. If the predictions come out as the theory says, then the theory is supported. If the predictions do not come out as the theory says, the theory is not supported and should be rejected. This process has the basic pattern that Popper described, but the idea that theories can be "supported" by observations is not a Popperian idea.

The term "support" is vague, but I think discussions of the hypothetico-deductive method generally assume that if a theory makes a lot of successful predictions, we have more reason to believe that the theory is true than we had before the successful predictions were made. We will never be completely sure, but the more tests a theory passes, the more confidence we can have in its truth. The idea that we can gradually increase our confidence that a theory is true is an idea that Popper rejected. As I said at the start of this chapter, some of Popper's scientific admirers do not realize that Popper's view has this feature, because some of Popper's discussions were misleading.

Other formulations of the hypothetico-deductive method include a first stage in which observations are collected and a conjecture is generated from these observations. Popper disagreed with this picture of scientific procedure because he argued that fact-gathering always takes place in a way guided by conjectures. But this is a fairly minor point.

Another term that some textbooks use in discussing scientific method is "strong inference." This term was introduced by a chemist named John Platt (1964). Strong inference is roughly a Popperian kind of testing together with a further assumption, which Popper rejected. This assumption is that we can write down all the possible theories that might be true in some area, and test them one by one. We find the true theory by eliminating the alternatives. For Popper, this is impossible; in any real case, there will be an infinite number of competing theories. So even if we eliminate ten or a hundred possibilities, the same infinite number still remains. According to Popper, all we can do is to choose one theory, test it, then choose another, and so on. We can never have confidence that we have eliminated all, or most, of the alternatives.

I have not discussed objections to Popper's theory of scientific change yet, but I will do so in the next few chapters.

What is Popper's single most important and enduring contribution to philosophy of science? I'd say it is his use of the idea of *riskiness* to describe the kind of contact that scientific theories have with observation. Popper was right to concentrate on the ideas of exposure and risk in his description of science. Science tries to formulate and handle ideas in such a way that they are exposed to falsification and modification via observation. Popper's formulation is valuable because it captures the idea that theories can appear to have lots of contact with observation when in fact they only have a kind of "pseudo-contact" with observation, because they avoid all risks. This idea is a real advance. Popper's analysis of how this exposure works has a lot of problems, but the basic idea is good.

4.6 *Further Comments on the Demarcation Problem*

Popper is on to something when he says that scientific theories should take risks. In this section I will try to develop this idea a bit differently. Popper was interested in distinguishing scientific theories from unscientific ones, and he wanted to use the idea of risk-taking to make the distinction. But this idea of risk-taking is better used as a way of distinguishing scientific from unscientific ways of handling ideas. And we should not expect a sharp distinction between the two.

The scientific way of handling an idea is to try to connect it with other ideas, to embed it in a larger conceptual structure, in a way that *exposes it to observation*. This "exposure" is not a matter of simple falsification; there are many ways in which exposure to observation can be used to modify and assess an idea. But if a hypothesis is handled in a way that keeps it apart from all the risks associated with observation, that is an unscientific handling of the idea.

So it is a mistake to try to work out whether theories such as Marxism or Freudianism are themselves "scientific" or not, as Popper did. A big idea like Marxism or Freudianism will have scientific and unscientific *versions*, because the main principles of the theory can be handled scientifically or unscientifically. Scientific versions of Marxism and Freudianism are produced when the main principles are connected with other ideas in a way that exposes these principles to testing. To scientifically handle the basic principles of Marxism is to try to work out what difference it would make to things we can observe if Marxist principles were true. To do this it is not necessary that we write down some single observation that, if we encounter it, will lead us to definitively reject the main principles of the theory. It will remain possible that a background assumption is at fault, and there is no simple recipe for adjudicating such decisions.

To continue with Popper's examples, Marxism holds that the driving force of human history is struggle between economic classes, guided by ongoing changes in economic organization. This struggle results in a predictable sequence of political changes, leading eventually to socialism. Freudianism holds that the normal development of a child includes a series of interactions and conflicts between unconscious aspects of the child's mind, where these processes include resolving sexual feelings toward their parents. Adventurous ideas like these can be handled scientifically or unscientifically. Over the last century, the Marxist view of history has been handled scientifically enough for it to have been disconfirmed. Too much has happened that seems to have little to do with class struggle; the ever-increasing political role of religious and cultural solidarity is an example (Huntington 1996). Capitalist societies have also adapted to economic tensions in ways that Marxist views about politics and economics do not predict. It remains possible to hang on to the main principles of Marxism, but fewer and fewer people handle the theory in that way anymore. Many still think that Marxism contains useful insights about economic matters, but the fundamental claims of the theory have not held up well.

Freudianism is another matter; the ideas are still popular in some circles, but not because of success under empirical testing. Instead, the theory seems to hang around because of its striking and intriguing character, and because of a subculture in fields such as psychotherapy and literary theory that guards the main ideas and preserves them despite their empirical problems. The theory is handled very unscientifically by those groups. Freud's theory is not taken seriously by most scientifically oriented psychology departments in research universities, but it is taking a while for this fact to filter out to other disciplines.

Sometimes people say that Freud's ideas were indeed scientifically successful, because the idea that our minds contain unconscious processes is alive and well in scientific psychology. In that very broad sense, yes, a part of Freud's view is still with us, and recent work in neuroscience has investigated the idea of unconscious thought processes in ways that are undeniably scientific (Dehaene 2014). But I don't think Freud should get much credit for that broader idea, and there is not much left in psychology of the little internal agents ("id," "superego," etc.) that Freud wanted to describe.

Evolution is another big idea that can be handled either scientifically or unscientifically. People (including Popper) have wondered from time to time whether evolutionary theory, or some specific version of it such as Darwinism, is testable. What observations would lead scientists to give up current versions of evolutionary theory? A one-line reply that biologists sometimes give to this question is "a Precambrian rabbit." J. B. S. Haldane, an important biologist of the early twentieth century, is often credited with the line. An evolutionary biology textbook by Douglas Futuyma expresses the same point more soberly: finding "incontrovertibly mammalian fossils in incontrovertibly pre-cambrian rocks" would "refute or cast serious doubt on evolution" (1998, 760). The one-liner is a start, but the real situation is more complicated. So let us look at the case.

The Precambrian ended around 540 million years ago (the term "Precambrian" covers a number of different periods in the history of the Earth, all before the Cambrian, which began about 540 million years ago). Suppose we found a well-preserved rabbit fossil in rocks 600 million years old. All our other evidence suggests that the only animals around then were soft-bodied invertebrates (many of them very strange indeed) and that mammals did not appear until over 300 million years later. Of course, a good deal of suspicion would be directed toward the finding itself. How sure are we that the rocks are that old? Might the rabbit fossil have been planted as a hoax? Remember the apparent fossil link between humans and apes that turned out to be a hoax, the "Piltdown Man" of 1908 (Feder 1996). Here we encounter another aspect of the problem of holism about testing—the challenging of observation reports, especially observation reports that are expressed in a way that presupposes other pieces of theoretical knowledge. But let us suppose that all agree the fossil is clearly a Precambrian rabbit.

This finding would not be an instant falsification of all of evolutionary theory, because evolutionary theory is now a diverse package of ideas, including abstract theoretical models as well as claims about the actual history of life on Earth. The theoretical models are intended to describe what various evolutionary mechanisms can do in principle. Claims of that kind are usually tested with mathematical analysis and computer simulation. Evolution on a small scale can also be observed directly in the lab, especially in bacteria, viruses, yeast, and some animals like fruit flies. The Precambrian rabbit would not affect those results. But a Precambrian rabbit fossil would show that somewhere in the package of central claims found in evolutionary biology textbooks, there are some serious errors. These would at least include errors about the overall history of life, about the kinds of processes through which a rabbit-like organism could evolve, and probably about the "family tree" of species on Earth. The challenge would be to work out where the errors lie, and that would require separating out and independently reassessing each of the ideas that make up the package. This reassessment could, in principle, result in the discarding of basic evolutionary beliefs, like the idea that humans evolved from nonhuman animals.

Over the past forty years or so, evolutionary theory has in fact been exposed to a huge and sustained empirical test, because of advances in molecular biology. Since the time of Darwin, biologists have been trying to work out the total "tree of life" linking all species on Earth, by comparing their similarities and differences and taking into account factors such as geographical distribution. The trees that were arrived at prior to the rise of molecular biology can be seen summarized in various picturesque old charts and posters.

More recently, molecular biology has made it possible to compare the DNA sequences of many species. Similarity in DNA is a good indicator of the closeness of evolutionary relationship. Claims about the evolutionary relationships between different species can be tested fairly directly by discovering how similar their DNA is and estimating how many years of independent evolution the different species have undergone since they last shared a common ancestor. As this work began, it was reasonable to wonder whether the wealth of new information about DNA would be compatible or incompatible with the family tree that had been worked out previously. Suppose the DNA differences between humans and chimps suggested that the human lineage split off from the lineage that led to chimps many hundreds of millions of years ago and that humans are instead closely genetically related to squid. This would have been an enormous shock for evolutionary theory, one of almost the same magnitude as the Precambrian rabbit.

As it happened, the DNA data suggest that humans and chimps diverged about six million years ago and that chimps, along with bonobos, are our nearest living relatives. Prior to the DNA data, it was unclear whether humans were more closely related to chimps or to gorillas, and the date for the chimp-human divergence was much less clear. That is how the grand test of our old pre-molecular family tree has tended to go. There have been lots of new discoveries, and a number of interesting adjustments to the old picture, but no huge surprises. The version of evolutionary theory that was worked out in the years before molecular genetics stood up pretty well.

Further Reading and Notes

Popper's most famous work is his book *The Logic of Scientific Discovery*, published in German in 1935 and in English in 1959. The book is mostly very readable. Chapters 1–5 and 10 are the key ones. For the issues in section 4.4 above, see chapter 5; for section 4.5, see chapter 10. A quick and clear introduction to Popper's ideas is the paper "Science: Conjectures and Refutations" in his collection *Conjectures and Refutations* (1963).

The Cambridge Companion to Popper (2016) is a collection of chapters reappraising Popper's ideas in different fields. It includes a paper of mine, "Popper's Philosophy of Science: Looking Ahead," that discusses some additional features of his view in a fairly positive light.

Newton-Smith's *The Rationality of Science* (1981) is an older book that is still useful for its clarity and its presentation of issues surrounding "corroboration." Salmon (1981) is an exceptionally good discussion of Popper's views on induction and prediction. Schilpp (1974) collects many critical essays on Popper, with Sir Karl's replies. For the story of Popper, Wittgenstein, and the brandished poker, see Edmonds and Eidinow (2001). Winther (2009) is a good discussion of prediction in evolutionary biology.

Popper's influence on biologists and his (often peculiar) ideas about evolution are discussed in Hull (1999). Horgan's book *The End of Science* (1996) contains a very entertaining interview with Popper. For a more informal exploration of the upsides and downsides of error in our beliefs, see Kathryn Schulz's *Being Wrong* (2010).