Chapter 10

Scientific Realism

10.1 *Strange Debates*

What does science try to describe? The world, of course. Which world is that? *Our* world, the world we all live in and interact with. Unless science has made some very surprising mistakes, the world we now live in is a world of electrons, chemical elements, and RNA molecules, among other things. Was the world of a thousand years ago a world of electrons and RNA? Yes, although nobody knew it back then.

But the concept of an electron is the *product* of debates and experiments that took place in a specific historical context. If someone said the word "electron" in 1000 CE, it would have meant nothing—or, at the very least, not what it means now. So how can we say that the world of 1000 CE was a world of electrons and RNA? We cannot; we must instead regard the existence of these things as dependent on our concepts, debates, and negotiations.

For some people, the claims made in the first paragraph above are so obvious that only a tremendously confused person could deny them. The world is one thing, and our ideas about it are another. For some others, the arguments in the second paragraph show that there is something badly wrong with the simple-looking claims in the first paragraph. The idea that our theories describe a real world that exists wholly independently of thought and perception is a mistake, and one linked to other mistakes—about the history of science, progress, and the trust and authority we should accord science today.

These problems have arisen several times in this book—in chapter 5 we looked at Kuhn's claim that when paradigms change, the world changes too, and in chapter 7 when we found Latour suggesting that nature is the "product" of the settlement of scientific controversy. I criticized those claims, but now it is time to have a close look at these problems.

There is a standard way of setting up the main debate in this area, which runs like this. We ask: does science tell us (at least sometimes, when things go well) about the hidden nature of a world that exists in a mind-independent way? *Scientific realists* say yes. Various opponents of scientific realism say no.

A "no" answer can be reached in a number of ways. Perhaps theories can never "reach beyond" experience in what they say. Perhaps theories can make claims that reach further, but we can't ever expect to get claims of that kind right. Or perhaps the world itself does not exist mind-independently. Views that oppose the "realist" picture are often quite different from one another. I am basically on the realist side in these debates. Certainly I oppose the opposition—all the views that are seen as the main rivals of scientific realism, I oppose. But I've come to think that many parts of the discussion are unsatisfactory, including a lot of what people say on the realist side. There is a collection of problems here, and it is best, I think, to address them one by one, and not to worry too much about whether the view you end up with is a version of "realism."

10.2 *Realism*

The term "realism" has no single meaning in philosophy. Sometimes realism about something—moral values, for example, or God—just means commitment to the existence of that controversial thing. Then the realism debate is a debate about God, or moral values, or whatever the particular topic might be.

There is also a tradition of debate about what our basic attitude should be toward the world that we seem to inhabit—the ordinary world of furniture and trees and so on. The familiar, common-sense view is that this world is out there around us, existing regardless of what we think about it. We are all parts of a single world that exists in space and time, learning about this world when things go well, and acting on it in various ways. That view is sometimes called "common-sense realism."

Despite perhaps being part of common sense, this view has been challenged repeatedly. One kind of challenge holds that we could never *know* anything about a world of that kind, so it makes no sense to think of the world of our everyday dealings as a world that exists mind-independently. We have to somehow bring the world of everyday life

"closer in" to the mind. In reply, you might try to show that we can learn about a mind-independent world, or argue that realism (in this sense) can be true even if knowledge of the world is limited or impossible; you might try to be a "skeptical realist."

The usual way to express common-sense realism is by saying that reality exists "independently" of thought and language, independently of what we think and say about it (Devitt 1997). I used that language above. But "independent" is too broad a term. People's thoughts and words are parts of the world, not extra things that somehow are outside or above it. When thoughts change, the world changes, because thoughts are part of the world. And thought and language have an important causal role in the world, by means of action. One of the main reasons for thinking, talking, and theorizing is to work out what to do, especially how to modify and transform things around us. Every bridge or light bulb is an example of this phenomenon. So it would not make sense to just say that realism asserts that "the world" exists "independently" of thought. But still, there is something worth saying here once we have made some qualifications. We could say it like this.

> *Common-sense Realism*: We all inhabit a common reality, which has a structure that exists independently of what people think and say about it, except insofar as reality is comprised of, or is causally affected by, thoughts, theories, and other symbols.

A realist, in this sense, accepts that we may all have different views about the world and different perspectives on it. Despite that, we are all here living in and interacting with the same world. The world I interact with is the world you interact with. We are both parts of the same world, and we can both affect it. This is a view in which the idea that a person's own perspective determines their own world is rejected. It's not true that "we each have our own reality," even though we might each have our own views about what is real.

I think that realism in this sense is right; the arguments against it are unconvincing. Let's now look at how these sorts of questions relate to the philosophy of science.

10.3 *Approaching Scientific Realism*

Suppose someone accepts a realist view in the sense described above. What role does science have? It would be easy and uncontroversial to add that science provides new ways for us to interact with various parts of the world, and it creates new tools and technologies. A realist will see the activity of science as part of the overall picture. The view called "scientific realism" is usually seen as saying more than this, though. One option is to see the scientific realist as asserting that the world *really is* the way it is described by our current and best-established scientific theories. We might say: there really are electrons, chemical elements, genes, and so on. The world as described by science is the real world—for the most part, anyway.

Others, such as Bas van Fraassen, argue that it is a mistake to express the scientific realist position in a way that depends on the accuracy of our current scientific theories. Scientific realism is a philosophical view about science as a whole that does not depend on what our current theories happen to say. If our current theories turn out to be false, we have made some scientific errors, but we might not have made any philosophical mistakes about how science works and what it means when things go well.

In response to this, the instinct of quite a lot of philosophers has been to qualify the claim made by a scientific realist—make it more cautious—but hang onto the basic idea above (Devitt 1997). We might say that scientific realism holds that *most of our mature scientific theories are at least approximately true*. "Mature" theories means those that have been carefully studied and survived a lot of testing. A mature theory in this sense might still be wrong, but then, realists say, its success would be a "coincidence" or a "miracle" (Smart 1963; Putnam 1975a). This is now called the "no miracles" argument—scientific realism is said to be the only view that does not make the success of science in prediction into a miracle. The kind of realism that the no-miracles argument is supposed to support is something like the claim that most of our mature scientific theories are at least approximately true.

Another option is to allow that a scientific realist can be much more pessimistic about whether we are getting things right, and perhaps whether we are likely to ever get things right. We might see scientific realism as a claim about the *aims* of science. There is a world we live in (as in common-sense realism), and science aims to describe it. That is basically how van Fraassen, who I mentioned above and will also return to later, thinks of the scientific realist position. Science might fail, but there is a real world whose structure we are trying to describe with science.

This is a "safer" kind of scientific realism, but it doesn't seem to be saying very much. People can have whatever aims they want. Is there such a thing as *the* aim of science, or even a list of aims? It is probably true that most scientists see themselves as trying to describe how the world really works, but that is just a matter of their attitudes.

We might then say that for scientific realism it's a *reasonable* aim of science to describe or represent the world as it really is. (That is how I did it in the first edition of this book.) The realist says there is a world our work is aimed at, and the tool kit with which we approach it—theories, mathematical models, observational testing, and so on—is one that could work if things go well. It could enable us to describe what is really going on.

That might sound trivial, but plenty of people have argued that there is a big obstacle, of some sort, to our describing how the world works. These people think the scientific tool kit is good enough for some purposes, but not for that one.

I will come back to those arguments in a moment, but first I want to make a broader comment about the relationships between a "realist" view of the world, on one side, and these questions about what attitude we should have toward theories that have been well tested, on the other.

One view that usually gets called a "realist" option is the idea that when a theory has been tested extensively, we have reason to believe what it says (most mature theories are at least approximately true). Something also usually seen as a realist position is the idea that we inhabit a world that exists mind-independently (in the qualified way spelled out in the previous section). However, there is a problem here, at least in principle. One of the things that science might tell us about is whether we live in a

world of that kind—a world that exists mind-independently—or not. In some interpretations, quantum mechanics shows us that we do not. According to some versions of this theory, the state of a physical system is partially determined by the act of measurement, in a still-unresolved sense of the term "measurement." A number of physicists have seen this as causing problems for common-sense realist ideas about the relation between human thought and physical reality.

Some people think this is just impossible; it can't be a scientific fact that there are no facts. This is one of the situations where the word "fact" causes problems, though (as it did in chapter 7—we will see a few of these). It can't be a true scientific sentence that there are no true sentences. That would make no sense. But it might be a true sentence, established by science, that the view of the world as containing mind-independent objects is wrong. That is, in principle, possible.

These interpretations of quantum mechanics are controversial, and my understanding is that they are fading in influence. Philosophers of physics tend to be more critical of these views than physicists themselves (e.g., Maudlin 2019). Even if those views are completely mistaken, that is not the point here. The point is that they show there is a possibility that science could conflict with common-sense realism, at least in some versions. So if we were to believe a theory of this kind—if we took a "realist" attitude to the theory—then we would have to drop "realism" in the other sense, the sense that involves mind independence. This is an illustration of problems with the whole idea of scientific realism as it is often understood. If realism combines a claim about the nature of the world we live in with a recommendation that we believe what science says, then there is the possibility of a clash between those two commitments.

I'll introduce another problem, too, this one about the idea that scientific realism involves a sort of overall confidence in the messages of science, such as the view that it would be a miracle if most of our well-tested theories were false. The problem is that different scientific fields might require very different treatments in this area. In a lot of discussions of scientific realism, people look for a general summary of how confident we should be in "science," or "current science," or "mature theories"—where this applies across all the different fields, from physics through

chemistry to biology and psychology and the social sciences, and also paleontology and meteorology and other interesting special cases. But we might have reason to have different levels of confidence, and also different *kinds* of confidence, in different parts of science. Ernan McMullin (1984) argued that we should not think of the parts of physics that deal with the ultimate structure of reality as a model for all of science. Physics is where we deal with the most inaccessible entities, those furthest from the domain our minds are adapted to dealing with. In physics we often find ourselves with powerful mathematical formalisms that are hard to interpret. These facts give us grounds for caution even when theories are empirically successful. Those special features of physics do not apply at all in the case of molecular biology. There we deal with entities that are far from the lowest levels, entities that we have a variety of kinds of access to. We do not find ourselves with powerful mathematical formalisms that are hard to interpret. Trying to work out the right attitude to have toward biology is not the same as trying to work out the right attitude toward theoretical physics.

In cases where theories are doing well, we might also have grounds for optimism about some features of a theory and not others. McMullin, John Worrall (1989), and others have developed versions of the idea that the confidence we should have about basic physics is confidence that some *structural* features of the world have been captured reliably by our models and equations. That is a special kind of confidence. There have been interesting cases in the past where a theory had the right structure even though it was, in many ways, quite wrong about the kinds of things that exist. Here is an example used by Laudan: Sadi Carnot thought that heat was a fluid, but he worked out some of the basic ideas of thermodynamics accurately despite this. The flow of a fluid was similar enough to patterns in the transfer of kinetic energy between molecules for his mistake not to matter too much. It's possible to believe that some parts of science are very good at describing structure—relationships between things—without being so good at describing the objects that stand in these relationships. This view, again, probably fits better with physics than other parts of science.

The issues in the last couple of pages should give some indication of

why I think that that standard ways of discussing scientific realism have problems. However, you might reply that the issues I talked about just above are a mostly a matter of fine-tuning a basically realist outlook. Perhaps so, and given this, it makes sense next to look at views that depart more dramatically from the picture I've been working with here. Indeed, many of the philosophers who have figured in earlier chapters of this book have seen part of their project as arguing for alternatives to scientific realism. As well as Kuhn and Latour, who I've mentioned already in this chapter, this can be said, in different ways, of Schlick, Goodman, Laudan, and others. So let's look at some rival views.

10.4 *Challenges from Empiricism*

Scientific realism has often been challenged by some forms of empiricism. One part of the debate about realism is often referred to as a debate between realism and empiricism, though I think this is not an inevitable opposition.

Traditional empiricists tend to worry about both common-sense and scientific realism, often for reasons having to do with knowledge. If there was a real world existing beyond our thoughts and sensations, how could we know anything about it? Empiricists believe that our senses provide us with our only source of factual knowledge, and many have thought that sensory evidence is not good enough for us to regard ourselves as having access to a world of the kind to which the realist is committed. And it seems strange (though not absurd, I think) to be in a position where you simultaneously say that a real world exists and also say we can never have any knowledge about it whatsoever.

In large part because of these debates about knowledge, various alternatives to realism were developed. "Idealism" is a term that can be used for several of them. *Subjective idealism*, as seen for example in the seventeenth-century empiricist George Berkeley, holds that the "physical" world is really comprised of a patterned flow of sensations, sent to

us by God. Later, various philosophers argued for similar views (minus God) by arguing that all we could *mean* when we talk about "the real world" is some sort of regularity in our experiences. That position is sometimes called *phenomenalism*. If this is right, then when we seem to make claims about external objects, all we are talking about is patterns in our sensations. Another form of idealism, which came more out of opposition to empiricism, holds that the world amounts to more than a flow of sensations, but it nonetheless has a mental or spiritual character as a whole—I discussed these views back in chapter 2.

Part of what logical positivism wanted to do was claim that all these debates are meaningless. But the logical positivist theory of how language works, in at least some versions, seemed to bring the view rather close to phenomenalism—remember Schlick, quoted earlier in this book, saying that "what every scientist seeks, and seeks alone, are . . . the rules which govern the connection of experiences, and by which alone they can be predicted." Even after these views faded, empiricist views of the meaning of scientific language stayed around, and tended to encourage the idea that all we can ultimately talk about is our own experience, so science can never hope to describe the structure of a world beyond our senses.

Views like that did not do well as explanations of how science works. Sometimes people say that different views of the working of scientific language won't be reflected in the practice of science, but I think this is not right at all. As J. J. C. Smart, an early defender of scientific realism, argued, if a scientist was really a phenomenalist—did not merely say this, but acted with this commitment—that person would be much more conservative than most scientists, in a particular way. Suppose all you really want to do is get rules for the prediction of experience, and you encounter a problem—your predictions run into trouble. Why not just make a small patch, a small change to your theory, to deal with the problem? Why should you suspect that you have gotten something seriously wrong? If you merely patch up your theory, you will probably make it more complicated and inelegant than it was before, but this does not seem a very big deal if all you want to do is predict events. So what if your predictive machine is not as neat and simple as you'd like? Smart

thought that a scientist who is a realist will treat problems differently—as raising the possibility, at least, that the theory's whole picture is wrong and should be reworked.

Empiricist views of scientific language, of the sort that cause problems with realism, do not fit well with how science works. But alternative views have been hard to develop. We are still waiting on a better theory of scientific language.

Not all empiricist critics of scientific realism have based their arguments in views of language. Bas van Fraassen, who has been an influential antirealist, confronts realism on the proper aims of science (1980). He suggests that *all we should ask* of theories is that they accurately describe the observable parts of the world. Theories that do this are "empirically adequate." An empirically adequate theory might also describe the hidden structure of reality, but whether or not it does so is of no interest to science. For van Fraassen, when a theory passes a lot of tests and becomes well established, the right attitude to have toward the theory is to "accept" it, in a special sense. (This is related to Laudan's sense of "acceptance," discussed in chapter 6, but not the same.) To accept a theory, for van Fraassen, is (1) to believe (provisionally) that the theory is empirically adequate, and (2) to use the concepts the theory provides when thinking about further problems and when trying to extend and refine the theory.

Regarding point 1, for a theory to be empirically adequate, it must describe *all* the observable phenomena that come within its domain, including those we have not yet investigated. Some of the familiar problems of induction and confirmation appear here. Regarding point 2, van Fraassen wants to recognize that scientists do come to "live inside" their theories; they make use of the theory's picture of the world when exploring new phenomena. But van Fraassen says a scientist can live inside a theory while remaining agnostic about whether the theory is true.

In discussions of scientific realism, the term "instrumentalism" is used to refer to a number of antirealist views. Sometimes it is used for traditional empiricist positions of the kind discussed above, but sometimes it is used in another way, one that I regard as more appropriate. Rather than saying that describing the real world is impossible, an instrumental-

ist will urge us not to worry about whether a theory is a true description of the world, or whether electrons "really exist." If we have a theory that gives us the right answers with respect to what we can observe, we might occasionally find ourselves wondering whether these right answers result from some deeper match between the theory and the world, but an instrumentalist thinks this question is not relevant to science.

Van Fraassen does not use the term "instrumentalist" to describe his view; he calls it *constructive empiricism*. The term "constructive" is used by so many people that it often seems to have no meaning at all, so I have reserved it for the views discussed below in section 10.5. I see van Fraassen's view as a kind of instrumentalism, but it does not matter much what we call it.

I think it is entirely possible for a scientist to aim only for empirical adequacy, in van Fraassen's sense. But he thinks that science should aim at no more than this. Why should we believe that? Why should science always stop at the observable? First, there is no sharp divide between the observable and unobservable parts of the world—unless you think that all you can observe is your own sensations. Some objects can be observed with the naked eye, like trees. Other things, like the smallest subatomic particles, can only have their presence inferred from their effects on the behavior of observable objects. But between the clear cases we have lots of unclear ones (Maxwell 1962). Is it observation if you use a telescope? How about a light microscope? An X-ray machine? An MRI scanner? An electron microscope?

Van Fraassen accepts that the distinction between the observable and the unobservable is vague, and he accepts that there is nothing unreal about the unobservable. He also accepts that we learn about this boundary from science itself. Still, he argues, science is only concerned with making true claims about the observable part of the world. But why should this be? Van Fraassen is saying that it's never reasonable for science to aim at describing the structure of the world beyond this particular boundary. Suppose we describe a slightly different boundary, one based on a concept a bit broader than observation. Let's say that something is *detectable* if it is observable or if its presence can be very reliably inferred from what is observable. As with van Fraassen's concept of observability, science itself tells us which things are detectable. In this sense, the

chemical structures of various important molecules such as sugars and DNA are detectable although not observable. So why shouldn't science aim at giving us accurate representations of the detectable features of the world as well as the observable features?

Perhaps our beliefs about the detectable structures are not as reliable as our beliefs about the observable structures. If so, we need to be more cautious when we take theories to be telling us what the detectable structure of the world is like. But that is fine; we often need to be cautious.

What is so special about the "detectable"? Nothing, of course. We could define an even broader category of objects and structures, which includes the detectable things plus those that can have their presence inferred from observations with moderate reliability. Why should science stop before trying to work out what lies beyond this boundary? We might need to be even more careful with our beliefs about those features of the world, but that, again, is no problem.

You can see how the argument is going. There is no boundary that marks the distinction between features of the world that science can reasonably aim to tell us about and features that science cannot reasonably aim to tell us about. As we learn about the world, we also learn more and more about which parts of the world we can expect to have reliable information about. And there is no reason why science should not try to describe all the aspects of the world that we can hope to gain reliable information about. As we move from one area to another, we must often adjust our level of confidence. Sometimes, especially in areas such as theoretical physics, which are fraught with strange puzzles, we might have reason to adopt an instrumentalist view, or care only about empirical adequacy in van Fraassen's sense. But that does not apply across the board; we can often hope to do more.

10.5 *Metaphysical Constructivism*

I'll now look at a quite different way of being opposed to realism. I will use the term "metaphysical constructivism" for a family of views including

those of Kuhn and Latour. These views hold that, in some sense, we have to regard the world as *created* or *constructed* by scientific theorizing. Kuhn expressed this claim by saying that when paradigms change, the world changes too. Latour expresses the view by saying that nature (the real world) is the product of the decisions made by scientists in the settlement of controversies. Nelson Goodman is another example; he argues that when we invent new languages and theories, we create new "worlds" as well (1978). For a metaphysical constructivist, it is not possible for a scientific theory to describe the world as it exists independent of thought, because reality itself is dependent on what people say and think.

These views are always hard to interpret, because they look so strange when we take them literally. How could we *make* a world (or the world) just by making up a theory? Maybe Kuhn, Latour, and Goodman are just using a metaphor of some kind? Perhaps. Kuhn sometimes expressed different views on the question, and he struggled to make his position clear. But when writers such as Goodman have been asked about this, they have generally insisted that their claims are not just metaphorical (1996, 145). They think there is something quite wrong with the realist picture. They accept that it's hard to describe a good alternative, but they think we should use the concept of "construction," or something like it, to express the relationship between theories and reality. The world as we know it and interact with it is partly (or entirely) a human construction. The way things are, or "the facts," is dependent on our beliefs, language, theories, or paradigms.

Some of these ideas can be seen as modified versions of the view of Immanuel Kant ([1781] 1998). Kant distinguished the "noumenal" world from the "phenomenal" world. The noumenal world is the world as it is in itself. This is a world that we are bound to believe in, but can never know anything about. The phenomenal world is the world as it appears to us. The phenomenal world is knowable, but it is partly our creation. It does not exist independently of the structure of our minds.

This kind of picture has often seemed appealing to philosophers who want to deny scientific realism, but do this in a moderate way. Hoyningen-Huene (1993) has argued that we should interpret Kuhn's views as similar to Kant's. In Michael Devitt's analysis of the realism

debates (1997), a wide range of philosophers are seen as either deliberately or inadvertently following the Kantian pattern. According to Devitt, constructivist antirealism works by combining the Kantian picture with a kind of relativism, with the idea that different people or communities create different "phenomenal worlds" via the imposition of their different concepts on experience. This relativist idea was not part of Kant's original view; for Kant, all humans apply the same basic conceptual framework and have no choice in the matter. The more recent, relativist views of this kind give rise to the idea that each language community, each culture, or perhaps even each person inhabits their own reality.

The Kantian picture is also seen as a way of holding on to the idea that there is a real world *constraining* what we believe, but doing so in a way that does not permit our knowing or representing this world. Though that move can be tempting, it is hard to see how we make any progress with these problems by adding an extra layer, "the phenomenal world" or something similar, between us and the real world that constrains us.

The term "social constructivism" is often used for roughly the same kind of view that I am calling metaphysical constructivism. But that term is used for more moderate ideas as well. If someone argues that we make or construct our theories, or our classifications of objects, that claim is not opposed to scientific realism. We do indeed "construct" our ideas and classifications; nature does not hand them to us on a platter. But realists insist that beyond ideas and theories there is also the rest of the world, much of which is not our product. In fields like sociology of science, however, there is an unfortunate tradition of not explicitly distinguishing between the construction of ideas and the construction of reality.

In history, philosophy, and sociology of science, I suggest that much of the motivation for recent metaphysical constructivism, and also for ways of writing that do not distinguish between the construction of ideas and the construction of reality, lies in a determination to *oppose* a particular kind of view, and to emphatically signal one's opposition to it. A lot of work in these fields has been organized around the desire to avoid a picture in which reality determines what scientists and other people think by stamping itself on the passive mind. The Bad View, the view to avoid, holds that reality acts on scientific belief with "unmediated

compulsory force" (Shapin 1982, 163). That picture is to be avoided at all costs; it is often seen as not only false but culturally harmful, because it suggests a passive, uncreative view of human thought. Many traditional philosophical theories are seen as implicitly committed to this Bad View. This is one source for descriptions of logical positivism as reactionary, helpful to oppressors, and so on. A result is a tendency for people to try to go as far as possible away from the Bad View. This encourages people to assert simple *reversals* of the Bad View's relationship between mind and world. Thus we reach the idea that theories construct reality.

Some people explicitly embrace the idea of an "inversion" of the traditional picture (Woolgar 1988, 65), while others leave things more ambiguous. But there is little pressure within the field to discourage people from going too far in these statements. Indeed, those who express more moderate denials of the Bad View leave themselves vulnerable to criticism from within the field. The result is a literature in which one error—the view that reality stamps itself on the passive mind—is exchanged for another error, the view that thought or theory constructs reality.

Putting all this together, where does this leave me with scientific realism? Well, here is a summary of some things I believe. I accept common-sense realism. This could be undermined by science—I don't think that is likely, but it is possible in principle. To think that common-sense realism could be undermined by science is to take scientific theories as more than tools for prediction. More generally, I think that scientific theories, when things go well, can in principle and often do in fact tell us about how the world works and what it contains, including some of its deeply hidden structure. We have to be cautious about the messages of the parts of science that deal with the most inaccessible and esoteric entities, but that doesn't mean we can't hope to get things right even here. On the other hand, I don't accept the no-miracles argument for scientific realism discussed in section 10.3; that is too simplistic a view of the ways in which theories can be empirically successful. I think that phenomenalism, metaphysical constructivism, and other standard views that oppose realism are false. When those contrasts are in mind, it's natural to see myself as on the realist side, but as I said earlier in this chapter, once a person's responses to all these different issues are worked out, it's not

very important to work out whether the view that results falls within the category of "scientific realism" or outside it.

10.6 *Underdetermination and Progress*

This chapter will end with two sections on particular topics in this area, one that was briefly discussed earlier and the other a little further away.

Above I said that scientific realism is often now seen (at least in large part) as a debate about whether we can be confident that we are getting things right—whether we should be optimistic or pessimistic about the aspirations of science to represent the world accurately. Some people argue that fundamental ideas have changed so often within science that we should always expect our current views to turn out to be wrong—just as many theories that people confidently believed in the past turned out to be wrong. Sometimes this argument is called the "pessimistic meta-induction." The prefix "meta" is misleading here, because the argument is not an induction about inductions; it's more like an induction about explanatory inferences. So let's call it "the pessimistic induction from the history of science." The pessimists give long lists of previously posited theoretical entities, like phlogiston and caloric, that we now think do not exist (Laudan 1981). Optimists reply with long lists of theoretical entities that once were questionable but which we now think definitely do exist, such as atoms, germs, and genes.

I said above that I think we should handle these questions in a field-by-field way; physics and biology have different sorts of histories and raise different questions. I'll now add some ideas on this topic.

Discussions of optimism and pessimism about science often take place against a backdrop provided by a general view, a claim called the "underdetermination of theory by evidence." In simple terms, this is the argument that for any collection of evidence we might have, there will

always be more than one theory that can, in principle, accommodate that evidence. If so, this seems to show that our preference for any particular theory must always be influenced to some extent by factors other than evidence—by simplicity, elegance, or sheer familiarity (see Psillos 1999, chap. 8, for a review). This idea is closely related to some holistic ideas about testing discussed earlier in this book, and Quine, again, is often associated with this argument. In chapter 2 we looked at the claim that when things go wrong in our predictions, there are always several different ways to adjust our web of belief to make sense of what happened. Now we are looking at a slightly different idea: at any time, given the data we now have (whether these data include surprises or not), there will always be more than one theory we could choose that fits with our data.

A lot of discussion of this theme has looked at a sort of extreme case, at the idea that there can be different theories that are compatible with *all* possible evidence. Maybe there are real cases of this, but many of the examples discussed are somewhat artificial ones involving deceptive demons or very minor differences between two theories. A different "underdetermination" argument, one that I see as more relevant to science, has been developed by Kyle Stanford (2006). He argues that there is a situation that science finds itself in over and over again, one with troubling consequences. When we look back on previous episodes in the history of science, a common situation is that the evidence people had back then, which they used to make their choice of a theory, was also consistent with another theory that they had not ruled out, because they had not even thought of that theory. We can say this because we now *believe* the theory that no one had thought of back then. If this is how things appear when we look back, why shouldn't we expect that future scientists will look back on us the same way? That is, why shouldn't we expect future scientists to say that we, today, should not have been so confident in our choices, because they (scientists in the future) will then be using a theory that we currently can't imagine? Stanford calls this the argument from "recurrent transient underdetermination." The argument is a sort of combination of the pessimistic induction from the history of science and the idea of underdetermination.

This will not be much of a problem if the usual case is one where the future theory is an improvement on earlier ideas, of a kind that shows the earlier ideas to have been on the right track. Most people working now would not mind being told later that their ideas were on the right track but needed improvement. The argument is troubling if the typical case is one where new theories tend to flatly reject the main ideas in old ones. That question can only be resolved by looking closely at the history of science, in different areas. We also need to keep in mind the possibility that a theory can appear, in retrospect, to have been quite misguided about what sorts of things the world contains but nonetheless have some structural features that we can recognize as being on the right track. A theory can get us close to some crucial relationships while being wrong in other ways (Worrall 1989).

There is also another way in which these underdetermination claims can be misleading. Stanford argues, as discussed above, that at each moment our data will be compatible with various different theories, including some we have not yet thought of. That is how things usually look in retrospect. But it's also compatible with this historical picture that whenever people find there are alternative theories that are compatible with all the current data, they go out and look for more data, and make a new choice between the theories. This will not be the end of the process, because the new collection of data will also be compatible with more than one theory. But then we can get even *more* data and try to discriminate between *those* alternatives. In this sequence of events, there is always underdetermination—we never get down to one theory. But there might seem to be a kind of progress, too. How significant this progress is might depend a lot on the relations between theories that we choose at different times—whether theories build on their predecessors, or just replace them. This affects whether we are really heading somewhere, or just traveling.

We could go further down this path, but I think there are some problems with the whole way the situation is being set up, and soon it will be time to correct them. That will happen in chapter 12, when we come to modern theories of evidence. Once we've examined those views, a lot of things in this section will look different.

10.7 *Natural Kinds*

To finish this chapter I'll look at a very different theme. In some discussions of realism, especially as it relates to science, people say that a realist position includes the view that the world is neatly divided up into "natural kinds," into categories (like *electron*, *gene*, and so on) that our scientific theories give labels to. Realism is said to include the idea that nature has definite "joints" or "seams," and an ideal scientific language would give us a perfect partitioning of the world—a categorization that God would agree with, if God wanted to do science. Sometimes the term "metaphysical realism" is used for this sort of position.

There is something right here. It would be an extremely minimal realism that just said there is *a* world we live in and deal with, perhaps one with no definite structure. But there are reasonable versions of realism that do not say that the world is prepackaged into neat kinds, or anything like that. And however this topic relates to realism, the problem of kinds and categories is interesting in its own right.

One of the main things language does is provide us with groupings of objects—words such as "human," "enzyme," and "automobile" do this. The term "nominalism" is sometimes used in philosophy for the idea that groupings are entirely imposed by us; any collection of things can be given a name, and the world does not compel us to categorize things in one way rather than another. An alternative view on this question is that some groupings are more natural than others—more in accordance with the way the world is. A view of this sort is often seen as both reasonable in its own right and potentially useful in solving other philosophical problems. Back in chapter 3, I mentioned attempts to solve Goodman's new riddle of induction with the idea that the word "green" picks out a natural kind while the word "grue" does not. More standard examples of natural kinds are the chemical elements—*hydrogen*, *helium*, and so on through the periodic table. (I use quotation marks when I am talking about words, like "hydrogen," and italics when I am talking about kinds or collections of things in the world, like *hydrogen*.) If one is looking for natural kinds, the chemical elements do seem to be good examples. To claim that they are natural kinds is not just to say that hydrogen atoms

are real, but that the kind *hydrogen* is real in a mind-independent way also—our word "hydrogen" picks out a collection of things that are unified by real similarity. In many sciences, the question of whether standard terms pick out real or natural kinds is the topic of ongoing debate, even if the philosophical terminology "natural kind" is not used. For example, are the mental disorders categorized in psychiatric reference books such as the *Diagnostic and Statistical Manual* (also known as the "DSM") natural kinds, or do we tend to apply labels like "schizophrenia" to a range of cases that have no real underlying similarity?

The example of schizophrenia also points us to another feature of the situation. In cases like this that involve the study of human beings, there are two varieties of dependence that a kind or grouping might be said to have upon us. First, there is the question of whether the category is arbitrary from a medical and biological point of view. Second, there is a separate question of what the effects are, on that collection of people, of having that categorization put in place. If a collection of people are categorized and described in a particular way, whether the grouping was initially arbitrary or not, this can have consequences of its own. These consequences can be mediated by both the treatment of those people by others (by medical workers, in the case of schizophrenia) and by those people's own self-conception.

Ian Hacking discusses this phenomenon, especially in relation to mental illnesses, in his book *The Social Construction of What?* (1999). Hacking distinguishes what he calls "interactive" kinds from "indifferent" kinds. *Schizophrenia* is an interactive kind; when the term "schizophrenic" is introduced and applied to people, it can lead to new behaviors in those people as they respond to this categorization. *Hydrogen* is an indifferent kind; it does not respond to our categorization in these ways. Hacking marks out interactive kinds in a narrow way; the members of the kind have to know and care about their classification. There is also a broader phenomenon here, in which the path through the world of a collection of objects (crop plants, for example, or an endangered species) is materially affected by our forming a new classification of them, and hence behaving differently toward them, even though the objects do not know it.

A problem with many discussions in this general area is the fact that there appears to be a lot of real *structure* in the world that is not orga-

nized into kinds. To pick a very simple example, differences in height between people are real (they do not depend on what we say and think), but tall people are not a natural kind. There is no border between the tall and the non-tall; there are just lots of gradations in height. Words like "tall," as opposed to "taller than," are almost always vague and indefinite in their application. The structure that exists in the realm of human height lends itself to mathematical description, and that, of course, is the usual way of describing height—in feet or centimeters. Part of the power of mathematics lies in its ability to represent structure of this sort.

Many of these questions about kinds, construction, interaction, and mind independence are brought together in a case that has been discussed extensively over recent years. This is the status of human racial categories, such as *black*, *white*, and so on. The discussion has gone through several phases. Back in the latter part of the twentieth century, attention to emerging knowledge in human genetics and an awareness of the artificiality of traditional racial categories led to a number of people arguing that races, as traditionally understood, do not exist at all (Lewontin 1972; Appiah 1994). The whole idea of races as *kinds* of people was seen as a mistake. This is sometimes called "eliminativism" about race—the traditional categories ought to be eliminated. More recently, this view has been seen as failing to recognize a kind of reality that races have as a consequence of human beliefs and institutions. "Constructivist" views of race hold that although human races are not biologically natural, they are nonetheless real, as a consequence of human politics, history, and attitudes. Being put into a racial category has consequences that are not in any sense illusory, though they do depend on human attitudes.

The views distinguished above are often presented as options that one should choose between—the literature asks us to choose between being a realist, a constructivist, or an eliminativist about race. But it seems better to recognize a number of different "kinds of kinds" here, with distinct roles. Michael Hardimon (2017) recognizes no fewer than four. In his framework, *racialist races* are groups of humans that are supposed to have different underlying biological natures that are associated with different abilities and moral characters. These racialist races, for Hardimon, do not exist; he is an eliminativist about racialist races. There is also a *minimalist* concept of race, in which people of a particular race merely

have some recognizable physical differences from others, indicating a particular geographical ancestry. Hardimon thinks that minimalist races are real, and also thinks this is the everyday concept of race. So he is not an eliminativist about these. *Populationist races* are collections of people with genetic similarities that reflect a particular ancestry; this is a biological analog of minimalist race. He thinks these races are real also. Lastly, there are what he calls *socialraces*. A socialrace is a collection of people who have been *taken* to comprise a racialist race (the first of his concepts above). Being a member of one of these groups can have many consequences, even though racialist races do not exist.

In a simpler framework, Adam Hochman (2017) uses the term "racialized group" for a group of people who have been treated as forming a race in the traditional and discredited conception. Those things—racialized groups—are real and consequential, even though the assumptions that formed the original basis for the categorizations are erroneous. These groupings are not "natural" kinds in a biological sense, but they are "natural" in not being merely arbitrary. The ways people are treated on the basis of racial assumptions and categorizations are genuine aspects of human social life and its history.

Much thinking about human races also tends to look for definite boundaries and kinds, even when there are gradations and mixtures. This is certainly true of the discredited racialist races in Hardimon's sense, but not only of those. In this respect, even Hardimon's minimalist races might in some cases be questionable.

The example of race is especially tangled, but it illustrates something general. In many contexts, a distinction between "natural" kinds and arbitrary or non-natural groupings is insufficient to capture what is going on.

Further Readings and Notes

Central works in the mid-twentieth-century resurgence of scientific realism include J. J. C. Smart's *Philosophy and Scientific Realism* (1963) and various papers collected in Hilary Putnam's *Mind, Language, and Reality*

(1975b). Leplin (1984) is a good collection on the debate as it developed in those years. In that book, Boyd (1984) gives an influential defense of scientific realism in the no-miracles tradition.

Devitt's *Realism and Truth* (1997) defends both common-sense and scientific realism. Psillos (1999) gives a very detailed treatment. Hacking (1983) and Fine (1984) are influential works on scientific realism that defend rather different views from those discussed here. Hacking focuses on the role of experimental intervention, while Fine tries to defuse the debate by criticizing both realist and antirealist commentaries on science.

On the puzzles surrounding quantum mechanics, Becker (2018) is helpful, along with the Maudlin book cited earlier (2019). I said that my understanding is that interpretations of quantum mechanics according to which the state of the world is dependent on the perceptions of scientists performing measurements are fading in influence, though as I was writing this chapter, a new paper defending this message appeared: Proietti et al. (2019). Schrödinger and Wigner took this view seriously. Einstein accused Bohr of a view like this, but Bohr apparently said Einstein had misunderstood him; see Schrödinger (1958) and Faye (2019). Here is Wigner:

> It was not possible to formulate the laws of quantum mechanics in a fully consistent way without reference to the consciousness. . . . It may be premature to believe that the present philosophy of quantum mechanics will remain a permanent feature of future physical theories; it will remain remarkable, in whatever way our future concepts may develop, that the very study of the external world led to the conclusion that the content of the consciousness is an ultimate reality. ([1961] 1995, 172)

One of the most popular recent approaches to the problem of natural kinds is the "homeostatic property cluster" view developed by Boyd (1991) and others. I don't think this approach helps very much, because the idea of "homeostasis" is merely a metaphor in most cases outside of some particular parts of biology. On the issue of gradients as opposed

to kinds in biology, see Gannett (2003). For additional discussions of race and racial kinds, see Spencer (2018a and 2018b), and Pigliucci and Kaplan (2003). For other discussions of properties and kinds, and their relevance to induction, see Armstrong (1983 and 1989), Dupré (1993), and Kornblith (1993).