

Chapter 4

Realism and anti-realism

There is a very ancient debate in philosophy between two opposing schools of thought called *realism* and *idealism*. Realism holds that the physical world exists independently of human thought and perception. Idealism denies this – it claims that the physical world is in some way dependent on the conscious activity of humans. To most people, realism seems more plausible than idealism. For realism fits well with the common-sense view that the facts about the world are ‘out there’ waiting to be discovered by us, but idealism does not. Indeed, at first glance idealism can sound plain silly. Since rocks and trees would presumably continue to exist even if the human race died out, in what sense is their existence dependent on human minds? In fact, the issue is a bit more subtle than this, and continues to be discussed by philosophers today.

Though the traditional realism/idealism issue belongs to an area of philosophy called *metaphysics*, it has actually got nothing in particular to do with science. Our concern in this chapter is with a more modern debate that is specifically about science, and is in some ways analogous to the traditional issue. The debate is between a position known as *scientific realism* and its converse, known as *anti-realism* or *instrumentalism*. From now on, we shall use the word ‘realism’ to mean scientific realism, and ‘realist’ to mean scientific realist.

Scientific realism and anti-realism

Like most philosophical ‘isms’, scientific realism comes in many different versions, so cannot be defined in a totally precise way. But the basic idea is straightforward. Realists hold that the aim of science is to provide a true description of the world. This may sound like a fairly innocuous doctrine. For surely no-one thinks science is aiming to produce a false description of the world. But that is not what anti-realists think. Rather, anti-realists hold that the aim of science is to provide a true description of a certain *part* of the world – the ‘observable’ part. As far as the ‘unobservable’ part of the world goes, it makes no odds whether what science says is true or not, according to anti-realists.

What exactly do anti-realists mean by the observable part of the world? They mean the everyday world of tables and chairs, trees and animals, test-tubes and Bunsen burners, thunderstorms and snow showers, and so on. Things such as these can be directly perceived by human beings – that is what it means to call them observable. Some branches of science deal exclusively with objects that are observable. An example is palaeontology, or the study of fossils. Fossils are readily observable – anyone with normally functioning eyesight can see them. But other sciences make claims about the unobservable region of reality. Physics is the obvious example. Physicists advance theories about atoms, electrons, quarks, leptons, and other strange particles, none of which can be observed in the normal sense of the word. Entities of this sort lie beyond the reach of the observational powers of humans.

With respect to sciences like palaeontology, realists and anti-realists do not disagree. Since fossils are observable, the realist thesis that science aims to truly describe the world and the anti-realist thesis that science aims to truly describe the observable world obviously coincide, as far as the study of fossils is concerned. But when it comes to sciences like physics, realists and anti-realists disagree. Realists say that when physicists put forward theories about

electrons and quarks, they are trying to provide a true description of the subatomic world, just as paleontologists are trying to provide a true description of the world of fossils. Anti-realists disagree: they see a fundamental difference between theories in subatomic physics and in palaeontology.

What do anti-realists think physicists *are* up to when they talk about unobservable entities? Typically they claim that these entities are merely convenient fictions, introduced by physicists in order to help predict observable phenomena. To illustrate, consider the kinetic theory of gases, which says that any volume of a gas contains a large number of very small entities in motion. These entities – molecules – are unobservable. From the kinetic theory we can deduce various consequences about the observable behaviour of gases, e.g. that heating a sample of gas will cause it to expand if the pressure remains constant, which can be verified experimentally. According to anti-realists, the only purpose of positing unobservable entities in the kinetic theory is to deduce consequences of this sort. Whether or not gases really *do* contain molecules in motion doesn't matter; the point of the kinetic theory is not to truly describe the hidden facts, but just to provide a convenient way of predicting observations. We can see why anti-realism is sometimes called 'instrumentalism' – it regards scientific theories as instruments for helping us predict observational phenomena, rather than as attempts to describe the underlying nature of reality.

Since the realism/anti-realism debate concerns the aim of science, one might think it could be resolved by simply asking the scientists themselves. Why not do a straw poll of scientists asking them about their aims? But this suggestion misses the point – it takes the expression 'the aim of science' too literally. When we ask what the aim of science is, we are not asking about the aims of individual scientists. Rather, we are asking how best to make sense of what scientists say and do – how to interpret the scientific enterprise. Realists think we should interpret all scientific theories as

attempted descriptions of reality; anti-realists think this interpretation is inappropriate for theories that talk about unobservable entities and processes. While it would certainly be interesting to discover scientists' own views on the realism/anti-realism debate, the issue is ultimately a philosophical one.

Much of the motivation for anti-realism stems from the belief that we cannot actually attain knowledge of the unobservable part of reality – it lies beyond human ken. On this view, the limits to scientific knowledge are set by our powers of observation. So science can give us knowledge of fossils, trees, and sugar crystals, but not of atoms, electrons, and quarks – for the latter are unobservable. This view is not altogether implausible. For no-one could seriously doubt the existence of fossils and trees, but the same is not true of atoms and electrons. As we saw in the last chapter, in the late 19th century many leading scientists did doubt the existence of atoms. Anyone who accepts such a view must obviously give some explanation of *why* scientists advance theories about unobservable entities, if scientific knowledge is limited to what can be observed. The explanation anti-realists give is that they are convenient fictions, designed to help predict the behaviour of things in the observable world.

Realists do not agree that scientific knowledge is limited by our powers of observation. On the contrary, they believe we already have substantial knowledge of unobservable reality. For there is every reason to believe that our best scientific theories are true, and our best scientific theories talk about unobservable entities. Consider, for example, the atomic theory of matter, which says that all matter is made up of atoms. The atomic theory is capable of explaining a great range of facts about the world. According to realists, that is good evidence that the theory is true. i.e. that matter really is made up of atoms that behave as the theory says. Of course the theory *might* be false, despite the apparent evidence in its favour, but so might any theory. Just because atoms are unobservable, that is no reason to interpret atomic theory as

anything other than an attempted description of reality – and a very successful one, in all likelihood.

Strictly we should distinguish two sorts of anti-realism. According to the first sort, talk of unobservable entities is not to be understood literally at all. So when a scientist puts forward a theory about electrons, for example, we should not take him to be asserting the existence of entities called 'electrons'. Rather, his talk of electrons is metaphorical. This form of anti-realism was popular in the first half of the 20th century, but few people advocate it today. It was motivated largely by a doctrine in the philosophy of language, according to which it is not possible to make meaningful assertions about things that cannot in principle be observed, a doctrine that few contemporary philosophers accept. The second sort of anti-realism accepts that talk of unobservable entities should be taken at face value: if a theory says that electrons are negatively charged, it is true if electrons do exist and are negatively charged, but false otherwise. But we will never know which, says the anti-realist. So the correct attitude towards the claims that scientists make about unobservable reality is one of total agnosticism. They are either true or false, but we are incapable of finding out which. Most modern anti-realism is of this second sort.

The 'no miracles' argument

Many theories that posit unobservable entities are *empirically successful* – they make excellent predictions about the behaviour of objects in the observable world. The kinetic theory of gases, mentioned above, is one example, and there are many others. Furthermore, such theories often have important technological applications. For example, laser technology is based on a theory about what happens when electrons in an atom go from higher to lower energy-states. And lasers work – they allow us to correct our vision, attack our enemies with guided missiles, and do much more besides. The theory that underpins laser technology is therefore highly empirically successful.

The empirical success of theories that posit unobservable entities is the basis of one of the strongest arguments for scientific realism, called the 'no miracles' argument. According to this argument, it would be an extraordinary coincidence if a theory that talks about electrons and atoms made accurate predictions about the observable world – unless electrons and atoms actually exist. If there are no atoms and electrons, what explains the theory's close fit with the observational data? Similarly, how do we explain the technological advances our theories have led to, unless by supposing that the theories in question are true? If atoms and electrons are just 'convenient fictions', as anti-realists maintain, then why do lasers work? On this view, being an anti-realist is akin to believing in miracles. Since it is obviously better not to believe in miracles if a non-miraculous alternative is available, we should be realists not anti-realists.

This argument is not intended to *prove* that realism is right and anti-realism wrong. Rather it is a plausibility argument – an inference to the best explanation. The phenomenon to be explained is the fact that many theories that postulate unobservable entities enjoy a high level of empirical success. The best explanation of this fact, say advocates of the 'no miracles' argument, is that the theories are true – the entities in question really exist, and behave just as the theories say. Unless we accept this explanation, the empirical success of our theories is an unexplained mystery.

Anti-realists have responded to the 'no miracles' argument in various ways. One response appeals to certain facts about the history of science. Historically, there are many cases of theories that we now believe to be false but that were empirically quite successful in their day. In a well-known article, the American philosopher of science Larry Laudan lists more than 30 such theories, drawn from a range of different scientific disciplines and eras. The phlogiston theory of combustion is one example. This theory, which was widely accepted until the end of the 18th century, held that when any object burns it releases a substance called 'phlogiston' into the

atmosphere. Modern chemistry teaches us that this is false: there is no such substance as phlogiston. Rather, burning occurs when things react with oxygen in the air. But despite the non-existence of phlogiston, the phlogiston theory was empirically quite successful: it fitted the observational data available at the time reasonably well.

find historical examples of empirically successful theories that later turned out to be false.

Whether these refinements can really save the 'no miracles' argument is debatable. They certainly reduce the number of historical counter-examples, but not to zero. One that remains is the wave theory of light, first put forward by Christian Huygens in 1690. According to this theory, light consists of wave-like vibrations in an invisible medium called the ether, which was supposed to permeate the whole universe. (The rival to the wave theory was the particle theory of light, favoured by Newton, which held that light consists of very small particles emitted by the light source.) The wave theory was not widely accepted until the French physicist Auguste Fresnel formulated a mathematical version of the theory in 1815, and used it to predict some surprising new optical phenomena. Optical experiments confirmed Fresnel's predictions, convincing many 19th-century scientists that the wave theory of light must be true. But modern physics tells us the theory is not true: there is no such thing as the ether, so light doesn't consist of vibrations in it. Again, we have an example of a false but empirically successful theory.

The important feature of this example is that it tells against even the modified version of the 'no miracles' argument. For Fresnel's theory *did* make novel predictions, so qualifies as empirically successful even relative to the stricter notion of empirical success. And it is hard to see how Fresnel's theory can be called 'approximately true', given that it was based around the idea of the ether, which does not exist. Whatever exactly it means for a theory to be approximately true, a necessary condition is surely that the entities the theory talks about really do exist. In short, Fresnel's theory was empirically successful even according to a strict understanding of this notion, but was not even approximately true.

The moral of the story, say anti-realists, is that we should not assume that modern scientific theories are even roughly on the right lines, just because they are so empirically successful.

XX
Examples of this sort suggest that the 'no miracles' argument for scientific realism is a bit too quick. Proponents of that argument regard the empirical success of today's scientific theories as evidence of their truth. But the history of science shows that empirically successful theories have often turned out to be false. So how do we know that the same fate will not befall today's theories? How do we know that the atomic theory of matter, for example, will not go the same way as the phlogiston theory? Once we pay due attention to the history of science, argue the anti-realists, we see that the inference from empirical success to theoretical truth is a very shaky one. The rational attitude towards the atomic theory is thus one of agnosticism – it may be true, or it may not. We just do not know, say the anti-realists.

This is a powerful counter to the 'no miracles' argument, but it is not completely decisive. Some realists have responded by modifying the argument slightly. According to the modified version, the empirical success of a theory is evidence that what the theory says about the unobservable world is approximately true, rather than precisely true. This weaker claim is less vulnerable to counter-examples from the history of science. It is also more modest: it allows the realist to admit that today's theories may not be correct down to every last detail, while still holding that they are broadly on the right lines. Another way of modifying the argument is by refining the notion of empirical success. Some realists hold that empirical success is not just a matter of fitting the known observational data, but rather allowing us to predict new observational phenomena that were previously unknown. Relative to this more stringent criterion of empirical success, it is less easy to

Whether the 'no miracles' argument is a good argument for scientific realism is therefore an open question. On the one hand, the argument is open to quite serious objections, as we have seen. On the other hand, there is something intuitively compelling about the argument. It really is hard to accept that atoms and electrons might not exist, when one considers the amazing success of theories that postulate these entities. But as the history of science shows, we should be very cautious about assuming that our current scientific theories are true, however well they fit the data. Many people have assumed that in the past and been proved wrong.


The observable/unobservable distinction

Central to the debate between realism and anti-realism is the distinction between things that are observable and things that are not. So far we have simply taken this distinction for granted – tables and chairs are observable, atoms and electrons are not. But in fact the distinction is quite philosophically problematic. Indeed, one of the main arguments for scientific realism says that it is not possible to draw the observable/unobservable distinction in a principled way.

Why should this be an argument for scientific realism? Because the coherence of anti-realism is crucially dependent on there being a clear distinction between the observable and the unobservable. Recall that anti-realists advocate a different attitude towards scientific claims, depending on whether they are about observable or unobservable parts of reality – we should remain agnostic about the truth of the latter, but not the former. Anti-realism thus presupposes that we can divide scientific claims into two sorts: those that are about observable entities and processes, and those that are not. If it turns out that this division cannot be made in a satisfactory way, then anti-realism is obviously in serious trouble, and realism wins by default. That is why scientific realists are often keen to emphasize the problems associated with the observable/unobservable distinction.

One such problem concerns the relation between observation and detection. Entities such as electrons are obviously not observable in the ordinary sense, but their presence can be detected using special pieces of apparatus called particle detectors. The simplest particle detector is the cloud chamber, which consists of a closed container filled with air that has been saturated with water-vapour (Figure 9). When charged particles such as electrons pass through the chamber, they collide with neutral atoms in the air, converting them into ions; water vapour condenses around these ions causing liquid droplets to form, which can be seen with the naked eye. We can follow the path of an electron through the cloud chamber by watching the tracks of these liquid droplets. Does this mean that electrons can be observed after all? Most philosophers would say no: cloud chambers allow us to detect electrons, not observe them directly. In much the same way, high-speed jets can be detected by the vapour trails they leave behind, but watching these trails is not observing the jet. But is it always clear how to distinguish observing from detecting? If not, then the anti-realist position could be in trouble.

In a well-known defence of scientific realism from the early 1960s, the American philosopher Grover Maxwell posed the following problem for the anti-realist. Consider the following sequence of events: looking at something with the naked eye, looking at something through a window, looking at something through a pair of strong glasses, looking at something through binoculars, looking at something through a low-powered microscope, looking at something through a high-powered microscope, and so on. Maxwell argued that these events lie on a smooth continuum. So how do we decide which count as observing and which not? Can a biologist observe micro-organisms with his high-powered microscope, or can he only detect their presence in the way that a physicist can detect the presence of electrons in a cloud chamber? If something can only be seen with the help of sophisticated scientific instruments, does it count as observable or unobservable? How sophisticated can the instrumentation be, before we have a case of detecting rather



9. One of the first photographs to show the tracks of subatomic particles in a cloud chamber. The picture was taken by the cloud chamber's inventor, English physicist C. T. R. Wilson, at the Cavendish Laboratory in Cambridge in 1911. The tracks are due to alpha particles emitted by a small amount of radium on the top of a metal tongue inserted into the cloud chamber. As an electrically charged particle moves through the water vapour in a cloud chamber, it ionizes the gas, and water drops condense on the ions, thus producing a track of droplets where the particle has passed.

than observing? There is no principled way of answering these questions, Maxwell argued, so the anti-realist's attempt to classify entities as either observable or unobservable is doomed to failure.

Maxwell's argument is bolstered by the fact that scientists themselves sometimes talk about 'observing' particles with the help of sophisticated bits of apparatus. In the philosophical literature, electrons are usually taken as paradigm examples of unobservable entities, but scientists are often perfectly happy to talk about 'observing' electrons using particle detectors. Of course, this does not prove that the philosophers are wrong and that electrons are observable after all, for the scientists' talk is probably best regarded as a *façon-de-parler*. Similarly, the fact that scientists talk about having 'experimental proof' of a theory does not mean that experiments can really prove theories to be true, as we saw in Chapter 2. Nonetheless, if there really is a philosophically important observable/unobservable distinction, as anti-realists maintain, it is odd that it corresponds so badly with the way scientists themselves speak.

Maxwell's arguments are powerful, but by no means completely decisive. Bas van Fraassen, a leading contemporary anti-realist, claims that Maxwell's arguments only show 'observable' to be a vague concept. A vague concept is one that has borderline cases – cases that neither clearly do nor clearly do not fall under it. 'Bald' is an obvious example. Since hair loss comes in degrees, there are many men of whom it's hard to say whether they are bald or not. But van Fraassen points out that vague concepts are perfectly usable, and can mark genuine distinctions in the world. (In fact, most concepts are vague to at least some extent.) No-one would argue that the distinction between bald and hirsute men is unreal or unimportant simply because 'bald' is vague. Certainly, if we attempt to draw a sharp dividing line between bald and hirsute men, it will be arbitrary. But since there are clear-cut cases of men who are bald and clear-cut cases of men who are not, the impossibility of drawing


a sharp dividing line doesn't matter. The concept is perfectly usable despite its vagueness.

Precisely the same applies to 'observable', according to van Fraassen. There are clear-cut cases of entities that can be observed, for example chairs, and clear-cut cases of entities that cannot, for example electrons. Maxwell's argument highlights the fact that there are also borderline cases, where we are unsure whether the entities in question can be observed or only detected. So if we try to draw a sharp dividing line between observable and unobservable entities, it will inevitably be somewhat arbitrary. But as with baldness, this does not show that the observable/unobservable distinction is somehow unreal or unimportant, for there are clear-cut cases on either side. So the vagueness of the term 'observable' is no embarrassment to the anti-realist, van Fraassen argues. It only sets an upper limit on the precision with which she can formulate her position.

How strong an argument is this? Van Fraassen is certainly right that the existence of borderline cases, and the consequent impossibility of drawing a sharp boundary without arbitrariness, does not show the observable/unobservable distinction to be unreal. To that extent, his argument against Maxwell succeeds. However, it is one thing to show that there is a real distinction between observable and unobservable entities, and another to show that the distinction is capable of bearing the philosophical weight that anti-realists wish to place on it. Recall that anti-realists advocate an attitude of complete agnosticism towards claims about the unobservable part of reality - we have no way of knowing whether they are true or not, they say. Even if we grant van Fraassen his point that there are clear cases of unobservable entities, and that that is enough for the anti-realist to be getting on with, the anti-realist still needs to provide an argument for thinking that knowledge of unobservable reality is impossible.

The underdetermination argument

One argument for anti-realism centres on the relationship between scientists' observational data and their theoretical claims. Anti-realists emphasize that the ultimate data to which scientific theories are responsible is always observational in character. (Many realists would agree with this claim.) To illustrate, consider again the kinetic theory of gases, which says that any sample of gas consists of molecules in motion. Since these molecules are unobservable, we obviously cannot test the theory by directly observing various samples of gas. Rather, we need to deduce from the theory some statement that can be directly tested, which will invariably be about observable entities. As we saw, the kinetic theory implies that a sample of gas will expand when heated, if the pressure remains constant. This statement can be directly tested, by observing the readings on the relevant pieces of apparatus in a laboratory (Figure 10). This example illustrates a general truth: observational data



10. Dialometer for measuring the change in volume of a gas as its temperature varies.

constitute the ultimate evidence for claims about unobservable entities.

Anti-realists then argue that the observational data 'underdetermine' the theories scientists put forward on their basis. What does this mean? It means that the data can in principle be explained by many different, mutually incompatible, theories. In the case of the kinetic theory, anti-realists will say that *one* possible explanation of the observational data is that gases contain large numbers of molecules in motion, as the kinetic theory says. But they will insist that there are other possible explanations too, which conflict with the kinetic theory. So according to anti-realists, scientific theories that posit unobservable entities are underdetermined by the observational data – there will always be a number of competing theories that can account for that data equally well.

It is easy to see why the underdetermination argument supports an anti-realist view of science. For if theories are always underdetermined by the observational data, how can we ever have reason to believe that a particular theory is true? Suppose a scientist advocates a given theory about unobservable entities, on the grounds that it can explain a large range of observational data. An anti-realist philosopher of science comes along, and argues that the data can in fact be accounted for by various alternative theories. If the anti-realist is correct, it follows that the scientist's confidence in her theory is misplaced. For what reason does the scientist have to choose the theory she does, rather than one of the alternatives? In such a situation, surely the scientist should admit that she has no idea which theory is true? Underdetermination leads naturally to the anti-realist conclusion that agnosticism is the correct attitude to take towards claims about the unobservable region of reality.

But is it actually true that a given set of observational data can always be explained by many different theories, as anti-realists maintain? Realists usually respond to the underdetermination

argument by insisting that this claim is true only in a trivial and uninteresting sense. In principle, there will always be more than one possible explanation of a given set of observations. But, say the realists, it does not follow that all of these possible explanations are as good as one another. Just because two theories can both account for our observational data does not mean that there is nothing to choose between them. For one of the theories might be simpler than the other, for example, or might explain the data in a more intuitively plausible way, or might postulate fewer hidden causes, and so on. Once we acknowledge that there are criteria for theory choice in addition to compatibility with the observational data, the problem of underdetermination disappears. Not all the possible explanations of our observational data are as good as one another. Even if the data that the kinetic theory explains can in principle be explained by alternative theories, it does not follow that these alternatives can explain as well as the kinetic theory does.

This response to the underdetermination argument is bolstered by the fact that there are relatively few real cases of underdetermination in the history of science. If the observational data can always be explained equally well by many different theories, as anti-realists maintain, surely we should expect to find scientists in near perpetual disagreement with one another? But that is not what we find. Indeed, when we inspect the historical record, the situation is almost exactly the reverse of what the underdetermination argument would lead us to expect. Far from scientists being faced with a large number of alternative explanations of their observational data, they often have difficulty finding even *one* theory that fits the data adequately. This lends support to the realist view that underdetermination is merely a philosopher's worry, with little relation to actual scientific practice.

Anti-realists are unlikely to be impressed by this response. After all, philosophical worries are still genuine ones, even if their practical implications are few. Philosophy may not change the world, but that doesn't mean it isn't important. And the suggestion that criteria

such as simplicity can be used to adjudicate between competing theories immediately invites the awkward question of why simpler theories should be thought more likely to be true; we touched on this issue in Chapter 2. Anti-realists typically grant that the problem of underdetermination can be eliminated in practice by using criteria such as simplicity to discriminate between competing explanations of our observational data. But they deny that such criteria are reliable indicators of the truth. Simpler theories may be more convenient to work with, but they are not intrinsically more probable than complex ones. So the underdetermination argument stands: there are always multiple explanations of our data, we have no way of knowing which is true, so knowledge of unobservable reality cannot be had.

However, the story does not end here; there is a further realist comeback. Realists accuse anti-realists of applying the underdetermination argument selectively. If the argument is applied consistently, it rules out not only knowledge of the unobservable world, but also knowledge of much of the observable world, say the realists. To understand why realists say this, notice that many things that are observable never actually get observed. For example, the vast majority of living organisms on the planet never get observed by humans, but they are clearly observable. Or think of an event such as a large meteorite hitting the earth. No-one has ever witnessed such an event, but it is clearly observable. It just so happens that no human was ever in the right place at the right time. Only a small fraction of what is observable actually gets observed.

XX The key point is this. Anti-realists claim that the unobservable part of reality lies beyond the limits of scientific knowledge. So they allow that we can have knowledge of objects and events that are observable but unobserved. But theories about unobserved objects and events are just as underdetermined by our data as are theories about unobservable ones. For example, suppose a scientist puts forward the hypothesis that a meteorite struck the moon in 1987.

He cites various pieces of observational data to support this hypothesis, e.g. that satellite pictures of the moon show a large crater that wasn't there before 1987. However, this data can in principle be explained by many alternative hypotheses – perhaps a volcanic eruption caused the crater, or an earthquake. Or perhaps the camera that took the satellite pictures was faulty, and there is no crater at all. So the scientist's hypothesis is underdetermined by the data, even though the hypothesis is about a perfectly observable event – a meteorite striking the moon. If we apply the underdetermination argument consistently, say realists, we are forced to conclude that we can only acquire knowledge of things that have actually been observed.

This conclusion is very implausible, and is not one that any philosopher of science would wish to accept. For much of what scientists tell us concerns things that have not been observed – think of ice ages, dinosaurs, continental drift, and the like. To say that knowledge of the unobserved is impossible is to say that most of what passes for scientific knowledge is not really knowledge at all. Of course, scientific realists do not accept this conclusion. Rather, they take it as evidence that the underdetermination argument must be wrong. Since science clearly does give us knowledge of the unobserved, despite the fact that theories about the unobserved are underdetermined by our data, it follows that underdetermination is no barrier to knowledge. So the fact that our theories about the unobservable are also underdetermined by our data does not mean that science cannot give us knowledge of the unobservable region of the world.

In effect, realists who argue this way are saying that the problem raised by the underdetermination argument is simply a sophisticated version of the problem of induction. To say that a theory is underdetermined by the data is to say that there are alternative theories that can account for the same data. But this is effectively just to say that the data do not entail the theory: the inference from the data to the theory is non-deductive. Whether the