CHAPTER 8

ANTI-REALISM

INTRODUCTION

In Chapter 7, we looked closely at the realist view of science. This takes the aim of science to be truth, not in some funny, post-modern sense, but in the sense of corresponding to states of affairs that are 'out there', in the world. And the main, some say 'ultimate', argument for this view is that realism is the only position that doesn't make the success of science a miracle. This is the 'No Miracles Argument', or NMA. In other words, just as theories are accepted – the realist claims – because they are the best explanations of the phenomena they are concerned with, so realism is the best (indeed the only) explanation of the success of science.

We then looked at the problems this position faces. First of all, the historically minded will say 'Been there, done that, and didn't like this view in the first place,' pointing out that throughout the history of science apparently successful theories have come and gone; theories that, had the realist been around at the time, she would have accepted as the truth, or close to it, but since they were subsequently thrown away as false, why should we believe that our current theories, amazingly successful as they are, should be regarded as true, or approximately so? This is known as the 'Pessimistic Meta-Induction' problem, or PMI.

Secondly, the situation may arise in which we have two theories that, it is claimed, are equally well supported by the evidence. This is the Underdetermination of Theory by Evidence problem, or UTE. If empirical success is supposed to be indicative of truth, how is the realist going to decide which theory is true, or closer to the truth? Now, the realist could always pin her hopes on further evidence coming in that will break the deadlock. But suppose that never happens? What could she appeal to then? She might suggest we go with the theory that is simplest, but then it's legitimate to ask what simplicity has to do with the truth. Or she could point to the way one of the theories is better integrated with our background knowledge than the other, but the same concern arises with regard to the background knowledge and so the problem is just pushed back a step.

And finally, there's the criticism that the NMA begs the very question at issue; that is, it assumes the very realist view it is designed to support. If you're not a realist, you're not going to accept the claim that scientists choose that theory which is the best explanation as true and so you're not going to be persuaded by the similar claim that realism is the best explanation of the success of science. Of course, the onus is now on you to come up with an alternative explanation of that success, but as we saw, that's not so difficult.

So, here's this chapter's fundamental question: How should we respond to these problems? There are various answers to this question out there in the literature but here I'm going to focus on just three, well-known and, I hope, interesting alternatives.

ALTERNATIVE 1: CONSTRUCTIVE EMPIRICISM

This is perhaps the dominant form of anti-realism around in the philosophy of science today. Basically, it identifies the source of the PMI and UTE problems as being the appeal to unobservable entities and processes and urges us to restrict our belief to observable things only. Now, it's important to be clear how this form of anti-realism differs from earlier forms, such as 'instrumentalism'. The instrumentalist, as the name suggests, took theories to be nothing but instruments for the prediction of empirical phenomena and as such, could not be regarded as true, or even approximately so. Theoretical statements – that is, statements about unobservable things like electrons, genes, the ego or whatever – are nothing but shorthand summaries of whole lists of observation statements. So, when a scientist states 'DNA is composed of a series of nitrogenous bases inter-connected by sugar and phosphate strands,' the instrumentalist takes this to mean 'When you do such an such an experiment, you will observe such and such a result.' (Obviously the list of observations in each case will be huge!)

The problem with this view is that it doesn't mesh with scientific practice. When a scientist says 'All the evidence is in and it looks as if our theory is pretty close to the truth,' the instrumentalist has to translate that as 'All the evidence is in and it looks as if our theory is a pretty good prediction device.' And when scientists talk about electrons, genes, the ego, whatever, the instrumentalist has to say 'Ah, what you're actually talking about are mammoth lists of observations' (to which the scientists might well respond 'No, what we're talking about are electrons, genes, the ego or whatever'!). On this view, we can't take the language of science literally and we have to translate all the talk and beliefs of scientists in terms of observations.

The constructive empiricist, on the other hand, takes the language of science literally. She agrees that when scientists talk about unobservable entities, their talk is, indeed, about these entities and is not mere shorthand for long lists of observation statements. And she also agrees that theories are the kinds of things that can be true. However – and here's the twist – the constructive empiricist adds what she takes to be a healthy dose of scepticism to the pot. How do we know that theories are true, she asks? In particular, how do we know that theoretical statements refer to the unobservable entities they purport to? If we buy into the empiricist premise that all knowledge is only of the empirical, that is, the *observable*, that is, what we can observe with the naked eye, then we clearly cannot *know* whether electrons, genes or the ego exist, nor can we know, therefore, whether theories are true or not. They might be, we just can't know.

On this view, then, we shouldn't believe theories to be true, or approximately so. What attitude should we take towards them? Well, what scientists do, as we have seen, is test their theories, seek empirical support for them, try to determine whether they are adequate in terms of accommodating the relevant observations. So, rather than believe theories as true, we should simply *accept* them as empirically adequate. As far as the constructive empiricist is concerned, this is the appropriate attitude we should adopt towards theories and, furthermore, we should drop the realist view that science aims at the truth and acknowledge that its aim is *empirical adequacy*. Here's what the 'founder' of constructive empiricism says:

> Science aims to give us theories which are empirically adequate; and acceptance of a theory involves as belief only that it is empirically adequate. . . . a theory is empirically adequate exactly if

> what it says about the observable things and events in this world are true exactly if it 'saves the phenomena'.[60]

What do theories tell us, then? On the realist view, theories tell us how the world is. But according to the constructive empiricist, we can never *know* how the world is, since we can never know its unobservable aspects. On this view, theories tell us how the world *could* be; that is, they provide us with useful stories of what the world might be like, but we can never know if these stories are actually true or not.

Now, you might find this a beguiling position; or you might think it's clearly mad! Before we make any critical judgments, let's see how it handles the three problems above.

So, first of all, how does constructive empiricism overcome the PMI problem? We recall that at the core of PMI lies the claim that the history of science presents us with case after case of radical change at the level of unobservable entities. Phlogiston, caloric, the ether have all been proposed by the respective theories and these theories have even enjoyed some empirical success, but they were all abandoned and these entities dismissed as unreal. Nevertheless, there is steady cumulative growth through the years at the level of the observable consequences of our theories. Of course, sometimes what initially appear to be good experiments are discovered to be flawed or problematic, but leaving these cases aside, the history of science does seem to present us with an accumulation of empirical results. (Some philosophers and sociologists have disputed this but we'll come to that in the next chapter.)

Now, it would seem that constructive empiricism can easily accommodate this. If we think of a theory as merely telling us how the world could be, then we shouldn't be surprised, or at all bothered, when the evidence tells us that no, it couldn't be like *that*. Of course, that doesn't provide conclusive evidence that the world *is* how the next theory proposes it is; again, this is just another way it could be. So the radical changes at the level of unobservables are nothing more than changes of story, from 'The world could be like this . . .' to 'Or it could be like that . . .'. In every case, we can't know for sure. And as the level of evidence accumulates, each successive theory can be viewed as more empirically adequate than its predecessor and so the growth of empirical knowledge can be accommodated.

Secondly, how does constructive empiricism overcome the UTE problem? This is even less of an issue. We recall that UTE states that there may arise situations in which we have two theories, both equally supported by the evidence, and hence we cannot believe either theory to be true. Indeed, says the constructive empiricist, nor should we! Nevertheless, we can accept both theories as empirically adequate. Of course, as a practising scientist, you may have to choose to work on one rather than the other simply because you don't have enough funding or expertise to work on both. Or you may decide that one is simpler, or easier to work with, than the other. That's all fine; the reasons for your choice have nothing to do with the truth of either theory. The choice to work with one or the other will be made on purely 'pragmatic' grounds.

Finally, what about explaining the success of science? Obviously, the constructive empiricist will not go with the NMA. Instead, she might question the sense in which, without the NMA, the success of science would be a miracle. As we noted at the end of the last chapter, this success appears striking but perhaps that is only because we focus on the successful theories and forget about all the others that fell by the wayside. Take a stroll through the university library sometime, run your fingers along the bound and collected volumes of the *Journal of Neurophysiology*, or *The Physical Review* or *The Journal of Chemical Ecology* or any one of the many specialist scientific journals, and pull out one of the dusty volumes from the early years. Just look at all the theories and hypotheses that were proposed but have subsequently been abandoned. Many of them, of course, were immature, half-baked even, but given such a plethora, is it any wonder that sometimes, some of them get it right?

Compare this, again, with the situation in biology: through mutations or recombinations, the DNA of an organism changes. These changes can be beneficial, harmful or neutral. If harmful, given a particular local environment, it may be unlikely that the offspring that inherit that mutation will themselves survive to reproduce and so the mutation dies out. If beneficial, again in the context of a particular environment, it may confer some advantage on the organisms that inherit it and hence the mutation spreads. As this process continues, entirely new species will form and we end up with the fox, which can live just about anywhere from the tundra to urban parks and can eat just about anything, and we think, 'Wow, This species is amazingly successful.' But it only seems miraculous if we forget about all the evolutionary false starts due to harmful mutations along the way. The actual explanation is rather prosaic: there were lots of changes, only a few of which were beneficial, and those are the ones we notice. It's the same with scientific theories: we tend to forget all the false starts and falsified hypotheses and by isolating the really successful ones treat that success as something in need of a realist explanation.

So, let's sum up the core ideas of constructive empiricism:

1) We have knowledge only of the *observable* (this is the empiricist feature), where what is observable is what could be observed by the naked eye in principle, that is, as described and understood by science itself. (So, for example, Jupiter's moons are observable, because science tells us that we could travel out beyond the asteroid belt and observe them with the naked eye; electrons are not observable because, despite certain cheesy science-fiction movies, science itself tells us that we could never shrink down to see them with our own eyes).
2) Unobservable entities and processes may exist but we can never *know*.
3) Theories may be true but we can never *know*.
4) Theories may nevertheless be accepted as empirically adequate.
5) Empirical adequacy, not truth, is the aim of science.

Now hopefully I've convinced you that this is an interesting view and, more than that, a viable alternative to realism. However, it too faces problems:

First of all, as we have just emphasised, it is grounded on the idea that we can have knowledge only of that which is observable with the naked eye. Now you might object that this relies on making a clear distinction between the observable and unobservable, one that we may have real trouble actually making. In times past, philosophers tried to draw the distinction in linguistic terms, between observation statements and theoretical ones, but gave it up as hopeless. The modern constructive empiricist doesn't think the distinction can be drawn that way, but rather in terms of the entities themselves. So, we're observable, and so are Jupiter's moons, but electrons are not. In between we might encounter a grey area where it's just not clear whether the entity concerned – very large molecules perhaps, or very small bugs – count as observable. But that just means that 'observable' is a vague term (like 'bald'), and as long as we have a good idea of when it can be employed and when not, there shouldn't be any problems.

More significantly, perhaps, you might feel that taking 'observable' to mean 'observable with the naked eye' is just way too restrictive. What about the use of instruments like the microscope? Don't scientists talk of 'observing' things through such instruments? Indeed, as we have noted, they even talk of 'directly observing' the core of the sun using highly specialised detectors that record the flux of subatomic particles known as neutrinos. However, this is where the constructive empiricist reminds us of that second term in her name – she is an empiricist and that means taking a particular stance with regard to what counts as knowledge, one that emphasises the role of experience, whether that be understood in terms of sensory data, or is extended to include the connections between these data. To insist that you have a broader understanding of 'experience' is just to adopt a different stance and if the constructive empiricist can account for everything you can, in particular if she can account for scientific practice, and, moreover, avoid the PMI and UTE problems, then it's not clear on what grounds you can say that your stance is better!

But what if we were to perform some Frankenstein type experiment and replace someone's eyes with twin electron microscopes? Such a person could presumably claim to 'observe' bacteria, the crystalline structure of various surfaces, even clusters of atoms (the 'Size and Scale' website (http://invsee.asu.edu/Modules/size&scale/unit3/unit3.htm) even categorises this under 'New Sets of Eyes'). Or imagine that the SETI project finally pays off and we find ourselves in contact with an alien species whose eyes have evolved differently so they can see to a level that we can't (just as birds, for example, can see polarised light that we can't). Doesn't that suggest that the constructive empiricists' distinction between what is observable and what is not is somewhat arbitrary?

The constructive empiricist answers as follows. We need to be clear that when we look at the beautiful images produced by a scanning tunnelling microscope, we are looking at images that have been produced by a physical process very, very different from the impact of light on the human eye, one that involves the 'tunnelling' of electrons between the surface of the body and a tip so sharp it consists of just a single atom, producing an electrical signal that is kept constant by raising and lowering this tip, and this raising and lowering is then recorded and enhanced by computer to yield an image.[61] So our Frankenstein operation would have to involve a bit more work than simply plucking out someone's eyes and replacing them with an instrument like this, and any alien that 'sees' via such a process would have to have a very different physiology from ours. Indeed, the constructive empiricist would insist, it would be so different that we would have to conclude that such aliens, or the results of our monstrous experiment, could not be counted as a part of our 'knowledge community', in the sense that what counts as knowledge for them would have to be very different from what it is for us.

The second major problem has to do with the explanation for the success of science. We saw that this provides the motivation for realism, through the NMA. Now is the Darwinian explanation of the success of science offered by the constructive empiricist itself adequate? Let's return to the analogy with evolution and 'survival of the fittest'. The underlying notion of the 'fitness' of a species is now understood in genetic terms, as we indicated in our sketch above. What would correspond to such terms for a *theory*? Let me put it another way: we now understand how a particular species, like the European fox, for example, is so successful in terms of the interaction between, ultimately, genetic changes and the particular environment the foxes' evolutionary ancestors found themselves in. The 'environment' for a theory might be taken as the empirical world, with experimental results leading to the extinction of certain theories and allowing the survival of others. But what would count as the underlying mechanism, analogous to the genetic make-up of an organism that drives the changes in theories? It's hard to see that there could be any such mechanism and so the analogy begins to look a bit ropey. The realist, of course, has an answer: a given theory is successful in that particular empirical environment because it has, in a sense, latched on to the world; it has 'got' the world right.

ALTERNATIVE 2: ENTITY REALISM

Even if you agree that constructive empiricism is too sceptical a stance and too restrictive in what counts as knowledge, you might be reluctant to return to full-blown realism. Is there no more modest form of realism that satisfies the feeling that we can know how the world is, in both observable and unobservable terms, and solves the problems faced by its bloated cousin? Here's another alternative that might do the trick.

First of all, let's recall that the source of the PMI problem is the apparent abandonment of certain unobservable entities throughout history; and the source of the UTE problem is the focus on the truth of theories. The view known as 'Entity Realism' (a view developed by Hacking and nicely described in his book, *Representing and Intervening*) offers a way through these difficulties by urging us to tear our philosophical attention away from theories and the thorny issue of whether they can be held as true or not, or might be true only we can never know it, or whatever, and instead focus on those unobservable entities that we are confident exist, not because they are presupposed by some theory, but because we *use* them. It's this pragmatic feature of entity realism that marks it out from other positions in the realism–anti-realism debate. The core ideas, then, are the following:

1) Some entities are *retained* through scientific change; e.g., the electron, the gene . . .
2) Our belief that these entities exist has nothing to do with the *truth* of theories, but with their *practical manipulation* in the creation of phenomena.

Here's what Hacking says:

> Experimental physics provides the strongest evidence for scientific realism. Entities that in principle cannot be observed are regularly manipulated to produce new phenomena and to investigate other aspects of nature. They are tools, instruments not for thinking but for doing. . . . The experimentalist does not believe in electrons because . . . they 'save the phenomena'. On the contrary, we believe in them because we use them to create new phenomena.[62]

He gives the example of scientists spraying a stream of electrons at a tiny niobium ball in order to change its charge in an experiment to detect the presence of sub-nuclear particles called quarks. We don't need to bother with the details of the experiment; what is important is the fact that the electrons are regarded as nothing more than just a tool which the scientists manipulate to create a new phenomenon. This has given rise to a famous slogan summarising Hacking's view: if you can spray 'em, they're real! The electrons are just something that can be more or less taken down off the shelf and used to achieve the desired effect. And just as the car mechanic doesn't fret over whether her wrench is real or not (at least not unless she's taken some philosophy classes), so the scientist doesn't and shouldn't worry about the reality of electrons and other unobservables.

Now, how does this position overcome the PMI problem? Well, the entity realist certainly acknowledges the issue and accepts that there is some change throughout the history of science at the level of unobservables; but the point of the PMI argument was to cut the realist's link between the empirical success of theories and belief in the existence of the entities posited by those theories. The theory that posited that heat was a kind of substance, called caloric, certainly enjoyed some significant empirical success but we now accept that heat is just molecular motion and there is no caloric. Nevertheless, there is also retention of certain entities at this level and thus we have grounds for *optimism*, but these grounds have nothing to do with the empirical success of the associated theories; rather they have to do with the use to which these entities are put. Consider the humble electron again: the associated theories have changed quite radically, from theories which took electrons to obey the classical mechanics of Newton, to the new quantum theory which suggested they had a wave-like aspect, to quantum electrodynamics which presents them as simply bumps in a quantum field, to today's string theories and so on. Despite all these changes, scientists have continued to believe in the existence of electrons because they have become an indispensable tool.

So, how does entity realism overcome the UTE problem? This is even easier to tackle. We recall that UTE insists that of two theories supported by the same evidence, we cannot believe either theory to be true, hence we cannot believe in either of the entities posited by the theories. Well, as we've seen, entity realism argues that belief in the existence of certain entities has nothing to do with belief in the truth of the associated theories. Indeed, Hacking maintains that scientists typically use different or even incompatible models of the electron, for example, without worrying about the *truth*. We can still believe that such entities *exist* even when faced with UTE-type situations.

Finally, then, how does this view explain the success of science? Remember, for the realist this is really important. She employs the same argumentative strategy as scientists themselves, or so she claims: namely, we take the best explanation of the phenomenon at hand and regard that as the truth. In the case of science, the phenomenon might be the bending of starlight around the sun, for example, and the theory would be Einstein's General Theory of Relativity; in the case of the philosophy of science, the 'phenomenon' is the success of science itself and the 'theory' is realism. However, the entity realist just isn't interested in the supposed truth of theories as this is not indicative for what we should take as 'real'. As we've already said, the empirical success of theories can be misleading, leading scientists to accept the existence of entities subsequently shown not to exist. The entity realist has a different view of success: science should be deemed successful not because it allows us to better represent the world, and say how the world is, but because it allows us to *intervene* in the world, by, for example, creating new phenomena and new technologies. It is intervention and not representation that we should be focusing on and it is the fact that we can use them as tools for intervention that leads us to believe in electrons and other unobservable entities.

Now, this is a powerful and quite complex view but it too faces certain problems.

First of all, it has an obviously unpalatable consequence; what if you're faced with an entity, or rather, with a hypothesis positing an entity, which you can't manipulate, and can't use to intervene in the world? The entity realist would presumably have to insist that you have no good grounds for taking this entity to exist. Now, stepping outside of the domain of physics for a moment, that may not present much of a problem for the chemist, say, since she can argue that as she uses certain kinds of molecules to produce certain effects and create certain kinds of phenomena, she can claim these molecules exist. Likewise, the biologist who uses certain enzymes to snip strands of RNA into pieces in order to create certain genetic phenomena has grounds for regarding these enzymes, at least, at real. But what about the psychologist who talks about the ego, say? They seem to be on much thinner ice. Perhaps that's a good thing, perhaps this is a way of winnowing out all the 'dodgy' entities and leaving only those we really should take as real (really real!).

But even in physics, or rather astrophysics, there may be problems. Astrophysicists have noticed a kind of phenomenon whereby very similar objects appear to be symmetrically reproduced across certain regions of space. Consider, for example, the 'Einstein Cross', where a 'cloverleaf' of four bright spots can be seen at the centre of a distant galaxy.[63] Now, most galaxies have only one nucleus, so this is an odd phenomenon and astronomers have tried to explain it by suggesting that in fact what we're seeing is the light from a far distant object known as a 'quasar', which is bent and split by the gravitational field of the intervening galaxy to produce four images. The galaxy is acting as a 'gravitational lens'.[64]

Now, many astronomers have come to accept the existence of these gravitational lenses because they explain a number of otherwise bizarre phenomena. And it's easy to see how the explanation goes: here's something really odd – four bright spots at the heart of a galaxy, for example. The chances that this is just a galaxy with a very unusual heart are really low; a better explanation – indeed, the best one – is that we're seeing another gravitational effect: the mass of the galaxy is so huge that it distorts the surrounding space-time sufficiently to refract and bend the light of a distant object, creating the four images. However, as far as the entity realist is concerned, this is not good enough; we cannot believe in the existence of gravitational lenses until we can use them and manipulate them to produce new phenomena. Now what are the chances that we're going to be able to use the centre of a galaxy the way a mechanic uses a wrench, any time in the future?! This puts the entity realist out of step with the top scientists in astrophysics, but perhaps that's a bullet she's prepared to bite.

That's not the only objection, however. The entity realist accepts that electrons, genes, etc. (but not gravitational lenses, or black holes), exist; but what are they? If we say, an electron is a charged sub-atomic particle, or a bump in a quantum field or the vibrating end of a quantum 'string' or whatever, where have we got that description from? A theory, of course. But how can we say what an electron, gene, or whatever, *is*, if our theories about it change, or if we have incompatible theories about it? As we've already noted, our description of the electron has shifted quite dramatically in the past hundred years or so, from being a small chunk of matter, to a wave-particle, to a bump in the Great Quantum Field, to the manifestation of a multi-dimensional superstring, to . . . But if we focus on these descriptions, we're faced with something like the return of PMI! The entity realist might be able to say that electrons exist, because she can use them like tools, but she can't say with any confidence what they are, because history teaches us that our current description may soon go the way of those held ten, fifty, a hundred years ago. However, if we can't say what the electron is, isn't our belief that it exists empty?

Again, the entity realist may just have to swallow that and agree that all she can say is that there's something out there, it's charged in such and such a way, it has the following mass, but that's all she can say. Here's a final objection that many people take to be a real hurdle that has to be overcome.

The entity realist, as we have seen, focuses on the use scientists make of certain entities. But are electrons, enzymes and the like really on a par with a mechanic's wrench? After all, you can't actually pull down a big box of electrons off the shelf and start throwing them around. What scientists actually do is use an electron gun, which produces an appropriately focused beam of electrons which can then be targeted on a niobium ball, or the inside of a TV screen or computer monitor. It's the electron gun that is more like a wrench, to be used by the scientists to achieve the effect they want. Now, the electrons are unobservable - that's the crucial issue, of course so what the construction and use of the electron gun rely on is an understanding of certain properties of the electrons (such as charge and mass) and the laws they obey. These laws may not be super-high level and abstract, they may be cobbled together in such a way that they apply only to the particular situations in which the electrons are produced, but it is these that the scientists rely on. In other words, the scientists must accept these low-level laws as true in order to achieve the effects they want. So, when we use electrons, say, to create new phenomena, we're relying on the truth of 'low-level' (causal) theories about electron behaviour. And these low-level laws and theories are accepted as true because they are empirically successful. But if we focus on these low-level laws, we're faced with something like the return of UTE! Suddenly entity realism doesn't look that different from the more standard form.

These sorts of objections haven't ruled the position out and many philosophers of science continue to develop it, particularly those who feel that the analysis of science tends to be too theory oriented and needs to focus more on pragmatic and experimental matters. However, there is another form of realism that goes to the other extreme and embraces the theoretical. Let's take a look at that before we move on.

ALTERNATIVE 3: STRUCTURAL REALISM

Let's go back to the PMI and look a little more closely at the history of science. And let's take another example: the history of light. Newton famously thought light was composed of tiny particles that underwent 'fits' when they passed from, say, air to glass, leading to the phenomenon of refraction. Then Young proposed that light is in fact a wave and Fresnel developed this theory further, producing a set of equations (now known as, surprise, surprise, the Fresnel Equations) which describe the behaviour of light when it passes from one medium – air, for example – to another – glass, say. We recall that when a critic pointed out that if light really was a wave, under the right conditions we should see a white dot in the shadow cast by an illuminated disc (due to diffraction around the edges of the disc), Fresnel ran the experiment and was as surprised as the critic when a white spot was observed. Maxwell brought light under the umbrella of his theory of electromagnetism (remember Hertz's experiments?) according to which it was conceived as an oscillating electromagnetic wave. Then came quantum theory and Einstein (again) argued that light had to be seen (!) as possessing particle-like qualities, so that it demonstrated the famous quantum wave–particle duality. Subsequently it too was regarded as a kind of quantum field, and so the story of its changing nature continues.

Now this looks like nice grist to the PMI mill: light as a Newtonian particle has been abandoned, as has light as a wave, so we have no good reason to suppose that in future years, the idea of light as a quantum field will also be consigned to the dustbin of history. But perhaps this is too hasty. Perhaps there is something that is retained throughout these dramatic theoretical shifts, something more than just all the empirical evidence that the constructive empiricist focuses on. After all, we still use Maxwell's equations (in certain circumstances) in the quantum era, and even after Maxwell proposed his theory, scientists still used Fresnel's equations. Indeed, they drop out of, or, more precisely, can be deduced from Maxwell's theory if certain conditions are applied, and in this sense they are retained despite all the shifts in our views of what light actually *is*. These equations can be understood as representing the underlying structure of reality and the view that *structure* is what is retained through theory change and is what we should be realists about, is known as *structural realism*. Its core ideas are as follows:

1) *Structure* is retained through scientific change.

2) *Structure* is what we should be realists about.

This is actually quite an old idea and if you look back through the last hundred years or so of commentaries on science, you find it cropping up again and again. Poincaré, for example, was a famous (and brilliant) mathematician and physicist (he came within a whisker of discovering the theory of relativity, for example), who also thought deeply about the nature of science. He also noted that certain equations are typically retained through theory change and wrote,

> . . . if the equations remain true, it is because the relations preserve their reality. They teach us now, as they did then, that there is such and such a relation between this thing and that; only the something which we then called motion, we now call electric current. But these are merely names of the images we substituted for the real objects which Nature will hide for ever from our eyes. The true relations between these real objects are the only reality we can attain.[65]

The idea then, is that all that we can know about reality is captured by the equations representing the relations between things, whose true 'natures' we can never really know (so to that extent the PMI is right).

How does this view overcome the PMI problem? Well, the answer should be obvious: PMI insists there is radical change at the level of unobservable entities; but it overlooks the fact that there is also retention of certain *structures* at this level. Shifting our attention away from the entities and on to the structures, it is the latter we should be realists about.

How does this view overcome the UTE problem? This is a little trickier, but one response runs as follows: UTE is supposed to lead us to conclude that we cannot believe either theory to be true, but that's OK, because the structural realist doesn't take the entire theory to be true, just those structural aspects that are retained through theory change. So, the structural realist will insist that in order for both theories to be empirically successful, they're going to have to possess certain equations or structures in common and it's that common part we should believe. Now, if the anti-realist can come up with examples of UTE where there are no common (structural) parts beyond the empirical level, then the structural realist will be scuppered. The extinction of the dinosaurs could be one example, although the structural realist will follow the ordinary realist in arguing that further evidence will surely settle that case one way or the other.

Finally, how does this view explain the success of science? Here the structural realist typically follows his non-structural cousin and argues that the success of science gives us good reason to suppose that our theories correctly describe the world, at least with regard to its structural aspects. In this sense the structural realist wants to present herself as less radical than the constructive empiricist and not as restrictive in her beliefs as the entity realist.

Now, as we noted above, this form of structural realism holds that all we can *know* is the structure of the world and we just have to remain agnostic about the nature of the entities. There is another strand, however, which insists that it is not that all we know is structure, but all there *is*, is structure. The motivation for this is quantum physics. A proponent of the above form of structural realism wrote that '[t]he structural realist simply asserts . . . that in view of the theory's enormous empirical success, the structure of the universe is (probably) something like quantum mechanical'.[66] But according to quantum physics, the 'nature' of the entities of the world as objects is deeply problematic. This is something the original heroes of the quantum revolution spotted and they noted that, according to the theory, the fundamental entities could not be regarded as individual objects, in the way that tables, chairs and people can. That's probably enough to make you wonder about the nature of these entities. It turns out, however, that the theory is consistent with the framework of individual objects. So now it appears we have another kind of fundamental underdetermination, only this time with the theory supporting two very different basic interpretations: in one the entities of the theory are individual objects, in the other they are not, in some sense. Anti-realists such as the constructive empiricist allege that this raises yet another problem for the 'standard' realist, since if she can't even say whether the objects she believes exist are individuals or not, what good is her realism?

This second form of structural realism responds to this challenge by suggesting that we should drop the notion of object from our theory altogether, so that what theories are about, on this view, are nothing but structures, pure and simple. Well, perhaps not so simple as it is not entirely clear what it could mean to say that the world is fundamentally just structure. A common understanding of a structure is that it consists of a family of relations holding over a set of objects. So, consider the genealogical structure of your family, with relations such as 'father of' and 'daughter of' holding between various people. But if the objects are removed from the picture, what do the relations hold between? And how can relations hold without any relata? These are crucial questions but to go any further would take us way over the cutting-edge and into the abyss! All I can say at this point is that explicating this form of structural realism is the top priority for these structural realists (such as the author!).

As with the other positions, structural realism also faces problems. First, in its focus on mathematical equations, this position seems to be oriented towards the more mathematical sciences, such as physics. What about biology, or even psychology, where there is much less mathematisation? Can structural realism find a place in these fields too? One answer is a blunt 'yes', since the notion of structure is broad enough that one can argue that maths is just one way of representing it. However, there is a lot more work to be done in developing structural realism in a biological context, say.

A second problem is associated with the question, is it always the case that structure is retained through theory change? What if the structures themselves change? If that happens, then we've lost one of the main advantages of going structural, which is to respond to the PMI. However, even granted that there has to be some change for science to progress, its not clear that the structures the realist is interested in change so radically that structural realism is fatally undermined. Finally, however, doesn't the above response to the UTE problem *assume* precisely what the realist needs to *show*? It merely expresses the hope that in any such cases, there will always be common structure. But what if two such empirically equivalent theories don't have any structure in common? Then we would see the return of UTE as well. As in the previous problem, what we need to see are some concrete examples, and these have not been forthcoming, at least not so far.

CONCLUSION

There are a variety of views on offer. The ones I've covered here – standard realism, constructive empiricism, entity realism and structural realism – are just some of the more well known. Which one you think is the 'best' account will depend not only on your understanding of scientific practice, its aims and its history but also on your philosophical views about what we can know. Any argument in favour of one position runs the risk of 'begging the question' against the others. All I've tried to do here is sketch the main arguments for and against, and bring you up close to the 'cutting-edge' in this area. Now we're going to consider a broader form of anti-realism, one that gets its force from the suggestion that scientific practice and, in particular, scientific change and progress, are not driven by observations and empirical support, but by social, political or economic factors.

Study Exercise 3: Truth and Existence

Consider the following questions:

Do you think bacteria exist? Why?/Why not?
Do you think genes exist? Why?/Why not?
Do you think electrons exist? (here it comes . . .) Why?/Why not?

Write your reasons down in each case. Are they the same *sorts* of reasons in each case or do they differ in some respect? Do you think any of these reasons are *better* than any of the others? If so, give reasons why you think they are better.

Now consider the following:

> There has been a huge discussion in recent years over the question of whether there is life on Mars. Not the evil aliens portrayed in *War of the Worlds*, but very simple life at the bacterial level. The discussion is important not simply because it's an interesting question whether life exists on other planets, but also because it has been suggested that life on Earth itself might be a by-product of life on Mars, as chunks of rock were blasted off the planet's surface by asteroid impact, carrying tiny hitch-hikers across the void.
>
> In 1996, NASA claimed to have discovered 'nano-fossils' – fossils of very small bacteria – in a meteorite that is known to have come from our neighbouring planet. A year later, NASA's Mars probe detected evidence that there had once been water flowing there, and if there was water, there had to have been an atmosphere as well. Further evidence has been produced showing channels shaped by water flow, as well as the possibility of large tracts

of frozen water still under the surface.

The claim was disputed, however, as other scientists argued that the 'nano-bacteria' would had to have been ten times smaller than any seen on Earth and that the so-called fossils were just remnants of chemical reactions within the rock. Then, five years ago, it was claimed that 'conclusive evidence' had been found of the remains of Martian 'magnetotactic bacteria'. This is a kind of bacteria which contains magnetic material called 'magnetite' and which was left behind when the bacteria decomposed within rock fissures, before being ejected from Mars by asteroid impact. The lead scientist, Dr Friedman, insisted that the crystals of magnetite his team had detected using electron microscopes, matched the criteria for such by-products of biology. He stated that, 'The chances of finding so many bacteria in a two kilogram rock from another planet are quite slim,' and 'It means bacteria must have been very widespread on Mars.'

Again, however, other scientists were not so sure. The UK's Colin Pillinger, head of a European Mars lander team, rejected the meteorite evidence as inconclusive and suggested that the only way to be sure was to land another probe on the planet, with the capability of detecting the chemical signatures of life. Unfortunately, his own lander crashed and burned in 2005 and so 'direct' evidence has still not been provided.

In 2006, two further arguments were put forward. Tiny cracks in another meteorite were found to be filled with a carbon-rich substance that resembles material found in fractures etched by microbes here on Earth. And very recently, a Russian-American team have suggested that the ability of certain bacteria to withstand very high doses of radiation is due to their having evolved on Mars, where radiation exposure can be a hundred times what it is here. Others, again, are not so sure, arguing that this ability (which has led to the nickname 'Conan the bacterium'!) is a mere side-effect of a defence mechanism against dehydration.

What sorts of reasons are being given here for the existence of life on Mars? Do they match the reasons you've given to the questions at the beginning of this exercise? Can you think of any grounds for doubting Dr Friedmann's conclusions? Why do you think Professor Pillinger's proposal would lead to more convincing support for the claim?
Do you think there is life on Mars?!