

Fundamental Physics and the Nature of Reality – Part 4

Philosophy and Sociology of Science

Norman Gray

October 2004

Preamble

Over the last three blocks, we have covered a dizzying range of topics in fundamental physics, and had our intuition repeatedly outraged.

With Special Relativity, we lost our confidence in time as a simple stream, and in spatial distance as an absolute notion that all observers agree on. With General Relativity and Cosmology we found out that space and time were not independent notions or things, and were forced to think of ourselves as living, not in the three spatial dimensions we're used to, but in a much richer four-dimensional spacetime which, through curvature, actually participates in the dynamics of the particles moving through it, rather than being the entirely passive background that Newtonian mechanics assumed it to be.

Quantum mechanics casts doubt on a huge range of our physical intuitions. Important parts of it seem to be non-deterministic; pairs of particles created together can seem to retain an identity even when they are widely separated; there seem to be limits in principle on the knowledge we can have of the world; there are doubts cast on the extent to which there *is* an external world. Quantum field theory (high-energy particle physics) supplements this with a description of Nature at very small scales. When it goes beyond this and attempts to describe Nature at unobservably small scales, attempts to discuss unified theories of *all* Nature which have some element of inevitability to them, and does so using very technical arguments relying heavily on abstract mathematical symmetries, (I claim that) it is placing a greater reliance on *a priori* reasoning than has been seen for several centuries in Science.

The theme that has been running through our consideration of these fundamental issues is: what is the picture of reality that modern physics has, and to what extent can we claim that this really is how the world is constructed? In the last few weeks we have gained some understanding of the conceptual background of modern physics. Now it is time to step a little way outside of physics, look back, and discuss just what it means to have a picture of reality, and how that picture is developed. I will start by giving an outline of the philosophy of science.

1 Philosophy of Science

Before we begin, we can ask the question: does science really need philosophy to justify it? The answer is no: there are many justifications of science – philosophical, technical, economic, aesthetic – which will appeal to a greater or lesser degree to different people. What the philosophy of science deals with is the nature of The Scientific Method (if such a thing actually exists), and the validity of science's implicit claim to a knowledge about the world which is more reliable, and even more profound, than that from any other system. That is, there are two different issues, here. Firstly, there is the question, which has historical and sociological aspects as well as philosophical ones, of how and why scientists choose between rival theories. Secondly, there is

the question of whether scientific knowledge – perhaps because of the answer to the first question – has any special status; whether it is in some sense ‘more right’ than knowledge of a different kind.

I will first describe three variant notions of *substance*, of what it is about a thing, that makes it that thing and no other. I will mention Descartes’ *rationalism* and the *empirical method* that was a partial response to it. I will finish off with accounts of the thought of *Karl Popper* and *Thomas Kuhn*, and their attempts to demarcate science and pseudo-science.

1.1 Substance

(...or, when is a cat not a cat?)

To what extent is a cat like other cats, and different from a tree? Why does an acorn grow into an oak and not, say, into an iguana?

Aristotle contrasted *substances* and *aggregates*. Aggregates are mere collections of matter, like a pile of sand, but substances are things which have a *real essence* – a set of properties which are necessary and sufficient conditions for an object to be a cat, or an oak, or whatever. For example, if an object has fur, pointy ears, four legs, a taste for fish, and so on, then it must be a cat, irrespective of what colour it is, and if it is missing any one of that set of properties, then it is not a cat. *Aristotle* elaborated this to include the notion of a mature form for each substance (a *telos*), and said that the natural evolution of the substance from the immature to the mature form (say, from a kitten to a cat, or an acorn to an oak) was a *necessary* change, and furthermore *explained* that change. Other changes that could take place, such as the kitten dying, or the acorn being eaten, were meaningless *accidental* changes.

Descartes adopted the notion of substances which have a real essence, but not the notion of a mature form for each substance. He furthermore suggested that there were only two substances: *mind*, the real essence of which was thought, and *matter*, the necessary and sufficient condition for which was spatial extension (that is, anything with spatial extension was matter, and all matter necessarily had spatial extent). Material things can change, but *Descartes* explains that each of these changes is the result of a mechanical interaction, rather than a teleological aspiration (towards a *telos*, that is).

In these terms, modern science also has substances, at many levels: it is a necessary and sufficient condition for an object to be a proton if it consists of two up and one down quarks stuck together; it is necessary and sufficient for an object to be a carbon atom if it has a nucleus with six protons in it; it is necessary and sufficient for a thing to be haemoglobin if it consists of so many carbon atoms, so many oxygen, so many nitrogen, . . . , in a particular arrangement. With *Descartes*, modern science will countenance only mechanical explanations as the real account of why things change.

1.2 Descartes, Hume and Newton

Descartes reached his position on the dualist distinction between mind and matter by a systematic consideration of what could and could not be doubted. He developed this into *rationalism*. In this view, the laws of nature must be generated by the mind working from self-evident first principles – from notions which cannot be doubted any more than the notion that two plus two equals four. For *Descartes*, the role of experiment and observation was to provide problems for the mind to solve, from first principles. This position necessarily gave mathematics a central role in science, and it led to *Descartes* denying the possibility that a vacuum could exist, and denying the possibility of action at a distance.

A much more fruitful philosophical position was the *empirical method*, which held that a theory of what a law of nature might be was only acceptable when confirmed by experiment or observation. However, ‘laws of nature’ are expressed in unobservable terms, in terms of ideal masses moving along frictionless planes, so they cannot be directly confirmed. This was refined into the *hypothetico-deductive method*: from the

abstract law, you deduce an observable prediction – ‘if you do such-and-such, then so-and-so will happen’ – if the prediction is in fact observed, then this is taken as *indirectly* demonstrating that the theory is in fact correct.

This was a tremendously powerful, and successful, way to think. Newton’s laws of motion were developed in this context and stood unchallenged through many demanding observational tests – the planets Uranus and Neptune were discovered exactly where Newton’s laws demanded they must be in order to account for irregularities in Saturn’s orbit. Neatly, it was an exactly analogous observational test – inexplicable irregularities in Mercury’s orbit – that demanded that Newton’s theory be eventually replaced by Einstein’s.

For all its practical success, however, *David Hume* (1711–1776) pointed out that the hypothetico-deductive method was logically flawed. It rests on the argument that ‘Theory *A* predicts observations *O*. Observations *O* are seen in fact, therefore theory *A* must be correct.’ It does not take account of the possibility that a completely different theory *B* might *also* predict observations *O* and furthermore correctly predicts as-yet unmade observations *P*, *Q* and *R*, which theory *A* gets wrong. That is, the method rests on the *inductive* supposition that ‘Theory *A* has got this (possibly large) set of observations correct. It is therefore true (and so will get all other observations correct)’.

Inductivism is the myth that The Scientific Method consists of scientists observing the world dispassionately and disinterestedly, and inducing from this mass of information some truth about the world. Inductivism, or at least the most naïve version of it, starts to crumble when we realise that we can never see the world except through some theory or other (are we observing a fire releasing phlogiston or a fire consuming oxygen?), and we can never even see some detail of the world without some background theory to highlight it (Michelson and Morley saw precisely nothing when they did their famous experiment, and it was only the fact that a detailed theory suggested that they perform that experiment and look for an aether drift, that gave that failure any significance).

1.3 Popper

Karl Popper (1900–1994) joined with Hume in deprecating inductive justifications for scientific truths – he agreed with Hume that no set of observations could lead you to rationally believe that a particular theory is true, or even probable. This led him to develop a philosophy of science – to describe both how science is done, and how it is justified – using the *deductive* techniques of *conjecture and refutation*. This is superficially similar to the hypothetico-deductive method, in that a theory will make an observational prediction, which is then tested and found to be true or false. However, while the hypothetico-deductive method would think of an experimental confirmation as a success, increasing the belief in the original theory, Popper would, apparently contrarily, regard the confirmation as a failure, as it has added nothing to our stock of knowledge (since we may not use the confirmation to inductively increase our belief in the theory). For Popper, a successful experiment is one which *fails*, and so allows us to reject the theory that produced the prediction. He doesn’t claim that scientists necessarily behave this way, but instead says that they are not being true scientists if they do not.

That final remark is important: Popper intended the principle of *falsifiability* to be a criterion of *demarcation* between science and pseudo-science (the ‘pseudo-sciences’ he had in mind seem to have been Marx’s view of history, and Freud’s, Adler’s and Jung’s views of psychoanalysis). He characterised as (real) scientists, those who will discard a theory out of hand when it is falsified by a contrary experiment. And he characterised as pseudo-scientists, those who, in the face of contradiction, will tinker with the details of a theory, or suggest that the experiment was done wrongly, or in other ways hang on to, or believe in, the theory in the face of adversity.

In this extreme form, the principle is far too high-minded to be reasonable – I cannot think of a single area of investigation which would qualify as real science under this criterion. In this straightforward picture, when Saturn’s orbit was found not to be an ellipse, Newton’s gravitational theory should have been abandoned. What happened instead was that scientists supposed that there was some factor in the observations that had not been taken account of in making the prediction of Saturn’s orbit; and a suitable extra planet was postulated, searched for, found, and named Uranus. This leads to a more sophisticated version of the criterion, under which a theory is rejected as pseudo-science if it invokes explanations for its predictive failure which are not, or which cannot be, supported by experiment for some reason.

This is still arguably not general enough. Even when the full set of planets was taken into account, there remained a tiny irregularity in the orbit of Mercury which Newton’s mechanics was quite unable to explain. There was no support for the existence of any other planets, so the problem was shelved, to await its resolution in General Relativity. The failure to reject Newton’s mechanics in the face of this experimental contradiction leaves us with a problem. We must either damn this failure as pseudo-science, and claim that any use of Newtonian mechanics after this observation was and is completely irrational and wrong, or else we must expand our demarcating criterion to allow a theory with minor observational problems still to be called ‘science’ if there is no reasonable theory around with which to replace it.

We are now in a quandry. If we are not to condemn as pseudo-science most of what scientists actually do (we could certainly do this, but it would leave us with an unhelpfully technical and specific definition of the word ‘science’), then we are obliged to expand the criterion to an extent which allows in areas which Popper, at least, certainly wished to exclude. Using this last version of the criterion, psychoanalysts and Marxist historians can claim (and have claimed) that their accounts may not be complete or even finally true, but that they should be accepted in the absence of anything better.

Popper approached the question from a logical point of view: following Hume, he wished to dispense with the need for any inductive justifications in science, and instead characterised science as an intellectual two-step, in which a hypothesis is *conjectured*, and various observational *predictions* are (deductively) obtained from it and from auxiliary ‘initial conditions’ (such as ‘there are nine planets’, ‘the planets’ mutual gravitational attraction is small enough that it may be ignored’). Then those predictions are tested by observation and if they are not found to be true, then the hypothesis, or theory, has been *refuted*, and should be abandoned. Popper described those who worked in this way as real scientists, and described as pseudo-scientists those who did not work in this way, but instead had the bad taste to believe in their theories, and made *ad hoc* adjustments to them, to explain failed experiments. This, he felt, was a clear distinction between science and pseudo-science.

Popper’s account *is* logical, it *is* free of induction, and it *is* a clear demarcation. However, it also bears very little relation to how science is actually done – there is (almost) nothing on the ‘science side’ of Popper’s demarcation. For example, if a prediction is falsified, it might make a lot more sense to disbelieve the initial conditions that allowed the prediction, than to disbelieve the theory behind them. It is generally difficult, if not impossible in principle, to falsify a fundamental theory. This is related to the theory-dependence of observation that I alluded to above: if no pure observation can exist, but only observations filtered through several meshes of theory, then the apparent falsification made by some particular observation could be due to the failure of any one of the theories involved in our perception of it.¹

Popper’s account can be elaborated, but as it is made more sophisticated, it loses much of its force.

We seem to be left without the clear demarcation that Popper wished. At best, it

¹For example, Copernicus’ cosmology was rebutted by the observation that Mars and Venus did not change in apparent size over the course of a year. This apparent falsification was however due to the wrong assumption that the naked eye could reliably measure the planets’ sizes, and not to an inadequacy of the theory the observation apparently contradicted.

seems to me, we are left with a set of criteria against which we can make judgements about *how scientific* an activity is – it becomes a matter of degree. Though Popper did not see his criterion as sociological, it does concentrate on the intellectual behaviour of individuals, so that our judgements should not be about particular incidents (“no, I will not reject this theory”), but instead be about a general willingness to reject or reconsider theories in the face of conflicting evidence.

What we can take from Popper are the notions that there is more force to falsification than to verification; and that critical evaluation of theories is important to science. Popper has raised some important issues, but he has not fully characterised the distinction between science and pseudo-science.

1.4 Kuhn

Thoman Kuhn provides an account of science which is radically different from Popper’s, both in its conclusions (that scientists normally simply accept the main body of their subject and are primarily concerned to fill in the details within that structure) and in its motivations (Popper moved from first principles to attempt to explain how science can be done without using induction, whilst Kuhn attempts to explain how science is done in fact).

Kuhn approached the problem from a historical, rather than a logical, point of view, and hoped to describe what science does and how it does it, by looking at what science has done in the past, and how it has changed. In *The Structure of Scientific Revolutions*, Kuhn claimed that the normal state of a science is for it to be working through the consequences of a particular *paradigm*. A paradigm is a ‘background theory’ that is overwhelmingly accepted by professional scientists, which informs the observations that they make and guides them towards further ones.

Kuhn described five stages which any science moves through: pre-paradigm, establishment, normal science, crisis and revolution.

1.4.1 Pre-paradigm

In the first stage, there are many candidates for a fundamental theory, and adherents of many schools. The discussion is about fundamentals, but there is little technical vocabulary, and so the discussion is generally intelligible. Because no-one has a generally acceptable theory of which things are fundamental, there is no agreement on which observations or measurements or questions are important, and which are irrelevant.

Greek cosmology was certainly pre-paradigmatic, with a large number of competing cosmologies floating around. It may be that psychoanalysis (and perhaps even psychic research) is a science-to-be (Popper would hate this!), in a pre-paradigmatic stage.

1.4.2 Establishment

At some stage, one theory will emerge from amongst its competitors and attract overwhelming support. This may happen because it has achieved continuing predictive success, or perhaps because the others collapse under the weight of inconsistencies and dead-ends. The other theories are not proved to be wrong, but are instead simply jettisoned and wither away as their adherents either switch to the consensus or die off. At this point, the science moves into the academy, becomes *the* fundamental theory, and is taught through increasingly technical textbooks.

This fundamental theory becomes the *paradigm* – the background theory which explains, and provides the words for, everything scientists see. The paradigm is now taken for granted, and is no longer routinely questioned, and generally no attempts are made to falsify it. You can even question the extent to which the paradigm *can* be falsified, if it provides the very words and concepts with which any falsifying experiment

would be expressed – that is, any falsifying experiments are themselves done within the paradigm.

Newton’s mechanics did precisely this in the late 17th century, with its talk of forces acting on particles which had inertia and mass.

1.4.3 Normal science

Once science has a fundamental theory which does seem to explain the world with some success, it can go about investigating and elaborating the paradigm while taking the fundamental theory for granted. In this stage, the question “what makes things fall?” is replaced by the question “how strong is the gravitational force acting on this mass?”, which is a question that could not have been asked before there was a consistent theory which talked of forces and masses. The discussion becomes more technical and specialised, and retreats into the pages of journals. Most science will be in this stage at any particular time. This is the domain of ‘professional scientists’.

Because scientists now have a clear picture of what is going on, they can ask very specific questions, with an expectation that the answers will be enlightening. That is, scientists can stop worrying about fundamentals, and instead ask detailed and technical questions, which might require a great deal of work to answer, because they can be sure that they will be rewarded with progress of some type. In contrast, a psychoanalyst (say) who spends years cataloguing the sock colours of patients with some particular neurosis, might be regarded as wasting his time, since there is no paradigm to suggest that whatever conclusions he comes to will have any meaning.

Thus there is no point in trying to measure the colour of an atom, since your theory of light and atomic structure tells you that colour is an attribute that atoms simply don’t have. However, there *is* a great deal of point in building an elaborate and possibly expensive apparatus to measure the atom’s mass and charge, since you can be confident that the atom would *have* a mass and charge and that their determination would have a bearing on other matters. You could estimate in advance, and explain in retrospect, your difficulties in finding them; and if the attempt is more difficult than anticipated, or even impossible, that is not a simple annoyance, but instead becomes a fact charged with great significance.

Accurate measurements will expose problems with theories within your paradigm, but will give hard information to allow you to adjust those theories in believable ways.

1.4.4 Crisis

Some of the problems that experiments expose will be easy to resolve. Some problems will be shelved, in the hope that other developments in the future will shed light on them; others will be accounted for by an unsatisfactory *ad hoc* theory, in the hope that further investigation, prompted by the temporary theory, will resolve the issue. However, there is a limit to how much this can be done. Eventually, problems and contradictions will accumulate unbearably, and the fundamental paradigm comes under attack.

In this stage, talk is again of fundamentals, and the discussion may even become generally intelligible again. Science again splits into camps, holding out for one resolution or another, or even for further modifications of the paradigm.

History is never as simple as historians would like it to be, but nonetheless we can identify some of the elements of such a Kuhnian crisis in physics at the end of the nineteenth century, with the problem of the aether, the continuing problems in understanding the spectrum of black-body radiation, and Boltzmann’s development of statistical mechanics.

Such a crisis can also be brought on by an unexpected discovery – when Darwin found the sub-populations of finches on the Galapagos islands, it started a train of

speculation which led directly to the evolutionary revolution (if that isn't too oxymoronic).

Artificial intelligence, where it overlaps with the philosophy of mind, is arguably in such a crisis now, as more elaborate and explicit theories of where consciousness comes from, and more elaborate machines which have 'intelligent' features, make the question "what is mind?" a very live one.

1.4.5 Revolution

From the mess of resolutions to the crises in the paradigm, one theory will emerge in the same way that the first paradigm emerged. It will absorb the old paradigm, it will additionally be seen to explain the crises which the old paradigm could not, and it will have enough richness to continue to throw up problems (that is, it will not simply 'explain away' the old crises – "...and then a miracle happens...").

The new theory will ask different questions from the old theory, in different language, and the preoccupations of the new theory will be as obscure to the old one as the old one's are to the new. It would be as senseless to us to try to weigh phlogiston as it would be to the pre-Copernicans to weigh a planet.

When the dust has settled, everyone has joined the consensus, and all the unbelievers have died or retired, a new paradigm will be in place, and the science will return to its 'normal' state. A big difference between Popper and Kuhn is that Popper would say that scientists discard a theory whenever it is falsified, but Kuhn acknowledges that theories hang on even when they are known to have problems. In a scientific revolution, it is the *success* of one of the replacement theories that drives out the old paradigm, and not the failure of the old paradigm by itself.

1.5 Lakatos and Feyerabend

Though I have perhaps represented Popper and Kuhn as opposites, only Popper (and possibly only early Popper at that) can really be said to be at one extreme. Between Popper and Kuhn lies *Imre Lakatos* and his notion of 'research programmes', and on the far side of Kuhn is *Paul Feyerabend* and his anarchistic notion that *all* forms of investigation and theory are ultimately of benefit.

1.5.1 Research programmes

Imre Lakatos presents a fundamentally Popperian picture, which can be described as 'sophisticated methodological falsificationism'. Lakatos talks of 'research programmes', which are mini-paradigms in the sense that they have a 'hard core' of background theory, surrounded by a 'protective belt' of auxiliary hypotheses. The core of the programme acts like a paradigm in suggesting experiments and predicting novel facts. When an anomaly is discovered, it is the protective belt that is modified or elaborated, and not the core, which survives as a guiding framework for the programme.

For Lakatos, a research programme is scientific if it places limits on the types of auxiliary modifications that may be made, and restricts them to testable or otherwise supportable ones. There may be many of these adjustments, but the process cannot go on for ever. Rather than being strictly falsified, Lakatos holds that a programme is rejected when it runs out of steam, and 'degenerates'. This will happen when the modifications have to be more arbitrary and more extreme, or when the hard core simply stops being able to suggest new directions.

Lakatos' picture is a more explicit version of one of the sophisticated variants of Popper's falsificationism that we have discussed already – it has a debt to Kuhn, though Lakatos repudiated Kuhn's 'socio-psychological' approach, and re-presented his insights 'through Popperian spectacles'. As falsificationism, however, it suffers from the same fundamental problem, that it is impossible to truly falsify a fundamental

theory. The notion of abandoning a degenerating programme is not enough, if you cannot give a clear account of just what makes a programme a degenerating one.

For example, I can claim that the continuing attempt to provide a hidden-variable interpretation of quantum mechanics is a degenerate programme, or even a completely stagnant one, as it does not seem to be generating any new insight or any experimental suggestions. But someone working in the area could quite reasonably say “we’re on the verge of a breakthrough – we’ll soon be able to produce discriminating experiments”, and I couldn’t prove them wrong. It comes down to the fact that I don’t *believe* that this programme will produce much in the way of interesting or useful knowledge.

As well as suggesting a research programme, a hard core of theory provides a rock on which a developing theory can become firmly established before being forced to confront observational anomalies. For example, Newton’s gravitational theory first demanded that the planets move on ellipses controlled by a fixed central sun, despite the fact that this contradicted his own third law of motion, which demanded that the gravitational force the sun exerted on the planets should be matched by a reactive force which the planets must exert on the sun, causing it to move in turn. Weaknesses in the theory manifested themselves as errors in prediction, and it was only Newton’s firm belief in the correctness of the law that made it worth while for him to develop the very sophisticated mathematical techniques needed to bring the reactive motion of the sun, and the influence of the planets on each other, into his calculations. That is, Newton *ignored* the observational anomalies (which we should regard as falsifications), as the existence of the mini-paradigm allowed him to distinguish these as secondary problems, distinct from the fundamental problems of the theory. The research programme which the paradigm suggested, gave him a map which guided him in developing the core of the theory until it became strong enough that it was able to address the anomalies constructively.

1.5.2 ‘Anything goes’

Paul Feyerabend supports Lakatos’ approach in general, but emphasises the problem of determining when a programme has degenerated. According to him, Lakatos *seems* to be describing a sophisticated falsificationism, but because there is no explicit test of whether a programme has degenerated or not, Feyerabend claims that this method is no method at all. Feyerabend takes this weakness of Lakatos’ theory, and turns it into the central pillar of his own: his resolution to the problem is to abandon the attempt to find a formal method of adjudicating between ideas, and instead allow an anarchistic free-for-all. In *Against Method*², he asserts that no scientific episode has ever been as simply intelligible as the philosophers of science have assumed, and that it has always been conventional scientific practice to ignore or explain away inconvenient facts, or to promote a particular theory or interpretation by rhetoric or propaganda. In the same book, he describes Galileo’s promotion of telescopic evidence in support of Copernicus’ cosmology, and describes the substantial differences between what he presumably saw through the telescope, what he said he saw, and what his critics and supporters saw in turn.

Feyerabend makes the historical points that our understanding of the world has always benefitted most from a zoo of ideas, and that “proliferation of theories is beneficial to science, while uniformity impairs its critical power”; and goes on to rejoice that conventional scientific approaches must compete on equal terms with mysticism and magic. Science is clearly more ‘successful’ than these other fields³, but for Feyerabend this is not enough to make it a qualitatively different *kind* of knowledge from magic, and he concludes from this the moral and social lessons that ‘[s]cience is

²Paul Feyerabend, *Against Method*, 1975. Third edition, Verso, 1993.

³... which rather begs the question, what does ‘successful’ mean? Possible answers might dwell on science’s superior abilities to fly, predict the future, and install itself as the established belief-system in the developed world. ‘Hegemony’ and ‘cultural imperialism’ are fashionable terms to use in this context.

neither a single tradition, nor the best tradition there is, except for people who have become accustomed to its presence, its benefits and its disadvantages. In a democracy it should be separated from the state just as churches are separated from the state'⁴, and talks of his "[a]nger at the wanton destruction of cultural achievements from which we all could have learned, [and] at the conceited assurance with which some intellectuals interfere with the lives of people. . ."⁵. For example, the emergence of complex ecological perspectives, the credibility of alternative medicine, and the appearance of drugs from the rainforest, are all phenomena which have become more or less respectable in recent years, but which were arguably antipathetic to the prevailing scientific culture a few decades ago. The fact that these topics have been (or are being) absorbed into new or expanding areas of science enhances the point, rather than diminishing it.

Apart from this final social point, Feyerabend's aim in *Against Method* is subversive, to undermine the notion that Science is a distinct form of knowledge, with a particular exclusive access to Truth, and that it gains that access by some identifiable procedure or attitude, which can be used to discriminate between legitimate and ignorable forms of knowledge. Whether or not you find Feyerabend's arguments persuasive, I believe he sets up a moral and philosophical context which we cannot shake off when we evaluate other approaches to science.

1.5.3 Theory choice and theory change

In talking about whether scientists do or do not subscribe to The Programme, Lakatos brings up the question of *theory choice*: how and why do scientists, individually or as a body, decide that one theory is correct and another is to be discarded. In a sense, this question captures what one arm of the philosophy of science is all about.

Lakatos sees the crucial choice taking place when a scientist decides to accept the core of a programme; Popper sees it happening when a scientist decides to accept a particular observation at face value; Kuhn, I think, sees the choice as being made by the community of scientists when they accept a theory as the revolutionary replacement for a now-discarded one.

Chalmers usefully distinguishes the problem of theory choice (why individual scientists choose one theory over another) from the question of theory change: how and why does one theory supplant another?⁶

1.6 Assessment

Kuhn's theory has some logical inevitability, as well as historical plausibility. Although there is some contact, there are substantial areas of dispute with Popper. Kuhn would discuss, and dismiss, pseudo-science, though he not completely subscribe to Popper's characterisation of it. The opposite is not true, however. Popper would dismiss Kuhn's mature, normal, science as pseudo-science, as it explicitly states that no attempts are made to falsify the background theory.

Kuhn's theory of science is *not* a sociological one. He does *not* claim, as some say he does, that each paradigm is no more than a fashion amongst scientific folk, and that there is no more truth to one consensus than to another. For Kuhn, each paradigm is an improvement on its predecessors, and science as a whole makes *progress towards truth*. For Kuhn, the paradigm is more than just a consensus (it is not a paradigm, for example, for everyone to believe in astrology), as it gives a direction to science, by suggesting profitable investigations, and allowing you to address yourself to certain

⁴*Ibid.*, Ch. 19, *Argument*.

⁵*Ibid.*, Ch. 20, *Argument*.

⁶A F Chalmers, *What is this thing called Science?*. Second edition, Open University Press, 1982. Chalmers suggests regarding 'degree of fertility' as an *objective* feature of a theory (that is, it is not a subjective judgement that any individual scientist will make, or even necessarily be aware of), and explains theory change in terms of the process of scientists 'colonising' the more numerous opportunities of a more fertile theory, and transforming the consensus by simple majority.

problems with the confidence that the solution will mean something, and increase your perception of the truth.

This is well expressed in a passage of Kuhn's which explicitly lays out his distinction between mature and pre-paradigmatic science:

First [of the criteria for a field's being a mature science] is Sir Karl's demarcation criterion without which no field is potentially a science: for some range of natural phenomena concrete predictions must emerge from the practice of the field. Second, for some interesting sub-class of phenomena, whatever passes for predictive success must be consistently achieved. (Ptolemaic astronomy always predicted planetary position with widely recognized limits of error. The companion astrological tradition could not, excepting for the tides and the average menstrual cycle, specify in advance which prediction would succeed and which would fail.) Third, predictive techniques must have roots in a theory which, however metaphysical, simultaneously justifies them, explains their limited success, and suggests means for their improvement in both precision and scope. Finally, the improvement of predictive technique must be a challenging task, demanding on occasions the very highest measure of talent and devotion. ('Reflections on my Critics', in I Lakatos and A Musgrave (1970) *Criticism and the Growth of Knowledge*, Cambridge)

Kuhn is also concerned with a different issue from Popper – that of distinguishing 'mature' science from pre-paradigmatic science. He gives as the criterion of a mature science that it has a paradigm, and claims that the distinction is important, as the existence of a paradigm means that a field will have directed and (scientifically) profitable research, is liable to lead to more or deeper knowledge, and is therefore worthy of a claim on resources.

In discussing Popper, Kuhn and Lakatos, and Feyerabend, we covered three quite different approaches to the philosophy of science. At the risk of caricaturing the complicated and varying positions they have held through their careers, we can describe Popper's approach as *nominalist* (describing how science ought to be done, irrespective of whether any scientists actually work that way), Kuhn's and Lakatos' as *descriptive* (attempting to justify how scientists' actual practice can lead to scientific knowledge), and Feyerabend's as *ethical* (fulminating against the unhealthy and stultifying authority science has arrogated to itself). Feyerabend appears to have moved through all three position in his career: if nothing else, I believe he illustrates how the philosophy of science is both a more substantial, and a more important, issue than we might at first expect.

2 Sociology of Science

Though Kuhn's description is part of the 'philosophy of science', in the discussion of normal science, and of the characteristics and effects of crises and revolutions, it is discussing the perceptions and beliefs of a community of 'professional scientists', and so obviously has some social aspects. We now bite the bullet, and briefly examine the claim that these 'social aspects', rather than being of no relevance, or at most being of merely methodological interest, are in fact of central importance, to the extent that scientific truth is *created* by the professional consensus, rather than being discovered.

From humble beginnings, science achieved a position of central authority in the West in the nineteenth and twentieth centuries, to the extent that it seems that only philosophy was deemed eligible to examine its workings and justifications. The erosion of this cultural and social dominance after the second world war emboldened sociologists to lift science's robes and examine its workings first as anthropologists, and later as critics.

These new approaches emphasise that there are substantial differences (as Feyerabend discussed in the case of Galileo) between what scientists think they do, and

what they appear to do in fact. By simply moving amongst the scientists, recording how they actually do their work (a claim which raises methodological problems of its very own), and examining the detailed histories of their published confrontations, sociologists were able to give an account of scientific life very different from the conventional one. This account rested much more heavily on rhetoric, authority and consensus than on the abstract scientific virtues of dispassionate investigation, uncontested deduction, and unassailable conclusion.

2.1 Scientific methodology

Collins and Pinch⁷ describe two such studies, which illustrate on the one hand the apparently unscientific causes of agreement on experimental results, and on the other the methodological problems which can inhibit such agreement.

First was the experimental test of Einstein's prediction of the bending of light-rays by passage near the sun. Einstein obtained the prediction through a derivation which some scientists had reservations about, but which Eddington, who was one of the few folk able to perform the calculation, did agree with. Eddington organised two expeditions to measure the actual deflection during the eclipse of 1919, one to West Africa and one to Brazil, and analyzed the experimental results himself. The measurements were particularly difficult to perform, and the numerical results were broadly scattered. However, of the three values obtained by the two expeditions, two (including the only 'good' value) worked out a little higher than Einstein's prediction, and the other was closer to the Newtonian prediction (obtained using Newton's gravitational theory, and his corpuscular theory of light): Eddington found reason to discard the lower figure, and did not mention that measurement when he wrote the (authoritative) account of the expedition. He used the two other measurements to declare Einstein's prediction confirmed; this implicitly confirmed the slightly disputed method of Einstein's calculation, and was explicitly taken as confirmation of General Relativity as a whole.

This account should not be taken to suggest that Eddington was in any way dishonest, nor even that the community should not have accepted the measurements as some verification; there is nothing in this account which is surprising to a practising scientist, even though it might not be quite the story she would tell when asked. Instead, this story illustrates a number of things about the normal progress of science. (i) There was no completely unambiguous prediction; instead, the results indirectly confirmed that Einstein's derivation was correct, and their reasonableness was in turn enhanced by the prediction. If the measurements had in fact failed to match either prediction, it is likely that it would have been those measurements, rather than the prediction, that would have been regarded as suspect. (ii) Motivated by his belief in the theory as much as by his knowledge of the details of the measurement, Eddington used his judgement to discard some readings and retain others, and his rhetorical sense in omitting those contrary readings from his description of the expedition. (iii) Eddington's high status, and the *imprimatur* of the Royal Society, where the results were announced, had a large part to play in the general acceptance of the results as correct, and as a full vindication of the General Theory.

A famous other example is provided by the lab logbooks of Millikan, doing his oil-drop experiment (subsequently used with great success to torture physics undergraduates, and break their spirit). It's a tricky experiment, and the logbooks show pages of runs, with one circled in red and a scrawled note 'publish this one'. Millikan knew perfectly well what the correct answer was – and he was right – because there was a very specific and fundamental theory which he was corroborating. A converse example is that of René Blondlot around 1903, and his consistent, but consistently unreproducible, measurements of 'N-rays'. Blondlot knew what he was looking for

⁷Harry Collins and Trevor Pinch, *The Golem: what everyone should know about science*. Cambridge, 1993.

when he did his measurements on N-rays, and repeatably and honestly saw it. He was wrong, however; N-rays didn't exist, and so this otherwise distinguished experimenter has gone down in history as a cautionary tale rather than a brilliant innovator.

A further point about the social difficulties of experimentation can be seen in the continuing attempts to detect gravitational waves, attempts in which the methodological problems have the additional clarity of being part of a current controversy, quite unclouded by any hindsight. GR indicates that violent movement of large masses, such as the formation or rotation of neutron stars, should excite gravitational radiation, which travels through spacetime at the speed of light and should be detectable on earth. The radiation is very weak and so, like Eddington's measurement of the deflection of light, the attempt to detect it is very difficult. Since the late sixties, many groups have built gravitational wave detectors, several have announced detection, and several have announced *non*-detection of waves (that is, they report that there are no waves above the minimum power their apparatus can detect). However, these announcements have all ended up being dismissed by those groups' rivals, on the grounds of inadequate equipment or inadequate experimental ability. The problem is that no one knows whether or not these waves actually do exist, and so whether or not anyone *should* be detecting anything at all. Collins has named this the *experimenter's regress*: we find out what the 'correct answer' is by looking at the results of competent experiments; an experiment can be judged to be competent if it obtains the 'correct answer'; since we don't know what the correct answer is, we can't decide which experiments are competent, and so indicate the correct answer. This cycle can only be broken by an independent consensus – that is one based on independent assessments of the quality of the apparatus and of the experimenters themselves – on which experiments are, and which are not, competent. It is difficult to see how the generation of this consensus can fail to have social or cultural aspects, or how it can be fitted into a satisfyingly logical philosophy of science.

As a final example, this time outside of physics, let us briefly consider biological taxonomy. Traditional, Linnean, taxonomy describes species of animals and plants based on the presence or absence of particular anatomical features; it is very successful in this, and was seen almost as a passive programme, uncovering the regularities within nature. Other biologists, concerned with whether particular organisms can cross-breed, and so share genes, classify species on this basis, and often come to different conclusions from the Linneans. Sometimes the Linneans will distinguish several species where the gene-swappers will see only one, and sometimes the gene-swappers will name several species which the Linneans regard as identical. Which group has the accurate picture of biological reality? I think this asks interesting questions about the reality of the objects we discuss in science, but the relevant point in this context is that two groups of scientists, with different scientific interests and motivations, have created two overlapping but non-identical versions of reality.

2.2 The sociology of scientific knowledge

These various approaches to the practice and justification of science leave us in a quandry: just what *is* the status of scientific knowledge? Although Kuhn was a 'convinced believer in scientific progress', and although a replacement paradigm usually *does* seem somehow 'better' than the theories it replaces, we must nevertheless acknowledge that there seems no *a priori* reason why that should be so.

If we cannot confidently say that one theory is 'closer to the truth' than another, then it becomes difficult to say just what scientific truth itself is. This is precisely the question that the philosophy of science tries to answer but, as we have seen, this has developed from Popper's logical certainties (falsifiability as the demarcation between science and non-science, and as the guarantor of scientific truth) to Kuhn's and Feynabend's social uncertainties. As it has done so, the description of science has come more and more to resemble what scientists actually do, but has simultaneously given

scientific culture a more and more prominent role in the formation of science's conclusions. In response to this, more recent, and more radical, sociology has moved from describing the culture which generates scientific knowledge, to discussing the social basis of that knowledge itself.

It is not particularly contentious to remark that scientific knowledge is underdetermined by the observations we make of the world – observations can fit in with numerous theories, in the way that GR and Newtonian gravitation *both* manage to explain most gravitational effects. Since scientists do in fact believe in certain theories and not in others, this leads us to ask what it is that creates this belief, and to suppose that at least some of the drive is social and cultural.

If scientific knowledge is partly conventional adequacy, and since these conventions vary from place to place and from time to time, we cannot say that the meanings of scientific concepts are stable, nor independent of circumstances. That is, scientific knowledge contains a large element that is *relative* to the social context in which it is generated.

There is a broad range of positions available within this radical programme, from a position not too far away from Kuhn or Feyerabend in suggesting that our notion of physical reality is more provisional than we might expect, all the way up to full ontological relativism, holding that all forms of knowledge are equally valid, and that our external reality is constructed by consensus and maintained by authority, rather than being discovered by objective experiment and observation.

One position in particular is that of the 'strong programme'. This programme investigates science whilst remaining impartial on the truth or falsity of its object, its rationality or irrationality, and crucially independently on whether the material it is studying is accepted into the mainstream or not (this is 'methodological relativism' – it deliberately discusses an episode without any consideration of whether or not the episode's conclusions turned out to be true in fact). That is, it studies both marginal *and* mainstream science, hoping to show that controversy and belief are part of *all* scientific activity, and that in a Kuhnian crisis, they are simply at their most obvious.

3 Assessment

So where does this leave us?

My own view is that science really is special in some way, in some way is more successful than magic or mysticism. There's a good argument to be had about just what 'special' and 'more successful' mean in this context, and whether they are or should be value-judgements, but I feel that if an argument concludes that science and magic are not distinguishable, then that argument has gone wrong somewhere, and the only interest it might yet provide is in finding out just where that error is.

I think that a large error in this programme is in confusing the processes which create knowledge with the processes which justify it: I may need a ladder to get onto the roof, but once there, it's not the ladder that supports me. The distinction, in the scientific case, is well made by Broad and Wade⁸, who talk of an 'invisible boot' of science. For Adam Smith, the 'invisible hand' of the community of free individuals – the free market – directs the development of that community independently of any guiding principle or controlling power. Similarly, the invisible boot 'kicks out all the incorrect, useless, or redundant data in science...[and] over time it stamps out the nonrational elements of the scientific process, all the human passions and prejudices that shaped the original findings, and leaves only a desiccated residue of knowledge, so distant from its human originators that it at last acquires the substance of objectivity'.

Even if we agree with none of these analyses of science, our understanding of what science is, how it produces new knowledge, and why it has its authority, are only enhanced by exposure to these radical ideas.

Copyright 1994–2002, 2004, Norman Gray.

⁸William Broad and Nicholas Wade, *Betrayers of the Truth*. Simon and Schuster, 1982. The authors study science through the 'pathology' of scientific fraud and self-deception.