Sarah Ormrod (ed) *Cambridge Contributions* (CUP, 1998) 122-142.

Cambridge Contributions: The Philosophy of Science

Peter Lipton

To admit at a cocktail party that one does philosophy of science is a good way to end the conversation. Many people have only the haziest idea what philosophers do and many people think that philosophy and science have nothing to do with each other. So I will begin with some general remarks about the philosophy of science, before turning to the great Cambridge tradition in the subject. Finally, because the only way properly to appreciate philosophy is to worry a philosophical problem for oneself, I will present a puzzle about the way scientists test their theories.

Justification and Description

What is the philosophy of science? The subject can be seen to emerge from more general areas of philosophy. One of the most important of these is epistemology, the theory of knowledge. The central issues are what knowledge is, how much of it we have, and how we acquire it. Epistemology often proceeds by presenting very negative, destructive arguments, arguments that seem to show that we do not know what we think we know, arguments that seem to show that we know almost nothing. Some of these sceptical arguments are familiar to adults and often reinvented by children. For example, you might worry that the distinctive experiences you have while, say, riding a bicycle, could in fact be experiences that you are having in the comfort of your bed. It could all be a dream that feels just like riding a bike, but none of it would be real. Given that these two situations -- the real bicycling situation and the dream situation -- feel, seem, exactly the same, on what rational basis do you believe the bike hypothesis over the dream hypothesis? That is the sort of argument that children and epistemologists worry about.

To understand better what knowledge is, how it works and how it changes, it helps to think about specific types of knowledge. If you made a list of the different types of knowledge we claim to have, scientific knowledge would probably come near the top. Science seems an example of knowledge acquisition at its most articulate, most ambitious and, many would say, most successful. So if one wants to understand knowledge, it pays to have a good look at science. That makes the connection to the philosophy of science, since much of the philosophy of science is the theory of knowledge with science as the example. We are trying to understand how scientific knowledge develops, how it changes. How does science work, how do scientific theories get produced, how are they tested, how are they evaluated, how do scientists weigh evidence? Those are the sorts of questions that one asks in the philosophy of science. They are not the only questions philosophers of science ask -- there are lots of others -- but they form a central part of the discipline.

Philosophical questions about scientific knowledge fall into two groups. In the first are the questions of justification. Are scientists really entitled to all the claims they make, or even to most of them? Are scientists are entitled to say that any of their theories are actually correct? Scientists often make ambitious claims, but the history of science is a graveyard of ambitious claims now rejected. So philosophers of science ask whether scientific methods can be justified and, if so, what those methods can be taken to produce. Do they produce the truth about the world, accurate predictions, reliable technology, helpful mythology, or what? For each of the aims science might have, one can ask whether the methods scientists use are suitable. Are the means suited to the ends? Can it be shown that the methods are really going to deliver what it is claimed they should deliver? These are questions of justifying science.

The second kind of question philosophers of science ask sounds more modest than the questions of justification. This is the question of description, the simple request for a general description of what is going on in science. Here the point is not to show what science really achieves or to defend its methods, but just to manage a better understanding of how science works, of what its methods are, for better or for worse. If this is all one wanted to do, just to describe how scientists test their theories, say, then one might think the job would not be very difficult. Just ask a friendly scientist and she will tell you.

Of course it does not work that way. There is a great gap between what people can do and what they can describe. It is one thing to be very good at riding a bicycle, quite another to be very good at giving a general account of how a bicycle is ridden, of the physics and the physiology involved. A person can be very good at doing it and very bad at describing it. To take another example, it is one thing to be able to speak a language fluently and so to be able to distinguish grammatical from ungrammatical strings in the language, but quite a different thing to be able to describe the principles that guide that judgement. Science is no different. Scientists may be very good at doing what they do, but they are not very good at describing what they do. I do not say this out of a feeling of philosophical superiority. Philosophers are pretty awful at describing what scientists do as well: it is just a very hard problem. But it is one of the central problems of the philosophy of science.

These are some of the questions philosophers of science ask; but what is the point? This is a question often asked about philosophy generally. Why isn't it all just a waste of time? What is the good of it? One of the reasons people ask this question is because philosophy does not seem to have any associated technology -- philosophy does not bake bread or build bridges. The philosophy of science, however, seems a possible exception, one of the few areas of philosophy where there might be a technology, broadly construed. Not bridges, but some people have hoped that the philosophy of science could make itself useful by helping scientists. That would be the practical application of the philosophy of science: it would make better scientists.

Unsurprisingly, many philosophers of science have been rather keen on this idea. Personally, however, I do not put much stock in it. The prospects of the philosophy of science providing extensive assistance to practising scientists nowadays are dim. Of course I do not conclude from this that the philosophy of science is a waste of time. An astronomer may devote his life to a better understanding of the stars without hoping to influence their behaviour. Similarly, a philosopher may hope to achieve a better understanding of how scientists work, without hoping to influence them. It may turn out to be useful as well as enjoyable for budding scientists to study some philosophy of science as undergraduates, but the justification of the discipline does not depend on this. Science is a central and pervasive part of our culture and our lives, and the attempt to understand better how it works and what it achieves is fascinating and worthwhile for its own sake.

The Cambridge Contribution

For the last forty years, there has been a thriving Department of History and Philosophy of Science at Cambridge, a Department that has become one of the outstanding centres for the history and philosophy of science in the world, with important work going on in many areas of the subject. Unparalleled library and archive resources and an extraordinary variety of research seminars attract scholars and graduate students from many countries. The Department also provides popular courses for undergraduate scientists, who can study history and philosophy of science alongside two other scientific subjects in their second year, and on its own in their final year. One of the secrets of our intellectual success is the productive way we combine the philosophy and the history, developing in tandem general accounts of how science works and particular accounts of how things actually went in particular places at particular times. We are also very fortunate to have one of the outstanding museums of the history of science and of scientific instrumentation in the world. Cambridge is a natural place to have such a wonderful resource, because Cambridge has been and continues to be the site of so much important and influential scientific research.

That is the last forty years in Cambridge for the history and philosophy of science - a world centre for the subject. But forty years is as nothing here; in forty years one is just beginning to figure out where the rooms are. To appreciate Cambridge's contribution to the philosophy of science, one has to go much further back. I propose to jump back not forty but four hundred years. This only takes us halfway, since the University is about eight hundred years old, but it already gives us an embarrassment of material, only a tiny part of which I can mention in this historical package holiday. Indeed I will only be able to touch on four Cambridge figures who have made important contributions to the philosophy of science; they will have to stand as representatives of many others.

Four hundred years takes us back to around the year 1600. The most famous figure in Cambridge philosophy of science at that time was Francis Bacon. He arrived at Trinity College at the age of twelve, and went on to become a lawyer, a politician, and one of the most influential philosophers of science of any period. Unsurprisingly, perhaps, his work in philosophy of science is more impressive than his work as a politician. Indeed as a politician he was corrupt, rightly accused of accepting bribes. He admitted as much, but offered the interesting defence that it should not be held against him since, although he took the money, he never let it influence him. This did not get him off, but he did not suffer too badly.

Whatever his political morality, Bacon was a gifted philosopher of science. He was a prolific and stylish writer, especially good at aphorisms such as the famous `knowledge is power'. Science has two components which Bacon sometimes referred to as `light' and `fruit'. The light that science provides is insight into the inner workings of the world, in particular, the workings of its invisible parts. But Bacon emphasised that science is not just about light, it is also about fruit, about technology, control, and improving the quality of life. That scientific knowledge should have this power is an obvious thought for us, but it may not have been nearly so natural for people in Bacon's time. The standard image of science then seems not to be nearly so practical, and what Bacon had to say on this subject may have played a role in changing that image.

Bacon also emphasised the fact that science is not static, but changes and grows and, he held, would progress if scientists handled themselves correctly. Bacon was one of those philosophers of science who thought that philosophers could help scientists to do better science. He emphasised the importance of careful observation, and the importance of gathering data without prejudice. Observation, he claimed, should come before the hypothesis it is supposed to test, lest the scientist's attachment to her hypothesis bias her observation. Once a hypothesis has been formulated, scientists should look not for data that might support it, but for negative instances, for counter-examples.

This influential idea is nowadays associated with the work of Karl Popper. Suppose the hypothesis (I am afraid philosophers' toy examples are dull) is that all ravens are black. No matter how many black ravens are observed to date, it remains possible that there is one of another colour lurking around the corner. In other words, positive instances will never prove a general hypothesis. One non-black raven is however sufficient to refute the hypothesis, because even if all the other ravens are black it is still false to say that all ravens are black. This is a striking logical asymmetry: positive instances never prove, but negative instances disprove. Bacon made a great deal of this. According to him, if scientists want to make scientific progress they had better spend a lot of their research time eliminating hypotheses by finding negative instances, in the hope that the hypotheses that survive will be true.

Much of Bacon's methodological advice was negative in a different sense. He held, rather plausibly, that people are prone to what is now called `systematic irrationality', and he set out to catalogue the various forms this irrational thinking could take, categorising them somewhat artificially under the headings of different `idols' of the mind. Thus Bacon reported that people tend to be overly impressed by evidence that confirms their prejudices, that they are misled in various ways by language, that they focus on superficial features of the objects they study, and many other forms of cognitive disability. One interesting feature of the diverse idols of mind is how widely they apply to thinking generally and not just to science. This fits with Bacon's view that the methods of science do not differ fundamentally from other forms of inquiry, including the day-to-day thinking we must do to manage our lives. This is a view for which I have considerable sympathy, though many philosophers of science, including some we will come to shortly, have been concerned rather to emphasise the differences between scientific and everyday thinking.

Like most philosophers, Bacon was much better at talking abstractly about what he supposed science was like than he was at actually doing any science. It is not that Bacon was uninterested in experimentation, but some of his experiments were rather odd. Indeed there is a famous story according to which his death was due to one of his stranger experiments. The experiment consisted of stuffing a chicken with snow to see if this was a particularly effective way to preserve the chicken. The fate of the chicken is not recorded, but Bacon is supposed to have died from the influenza he contracted while conducting the experiment. In addition to his dubious experimental technique, Bacon turns out not to have been the most perceptive judge of the science of his time. He seems, for example, to have ignored the work of William Harvey, the man who is credited with discovering the circulation of the blood, something we now regard as one of the major episodes in the history of science. What is strange is not just that Bacon gave no credit to Harvey, but that he appears not to have been aware of what Harvey was doing. This is particularly surprising since, as it happens, William Harvey was Bacon's personal physician.

If we now jump forward a century, we come to the Cambridge figure everyone has heard of: Isaac Newton. Another Trinity man, Newton is perhaps the greatest figure in the history of Cambridge and one of the greatest scientists that has ever lived. He gave us a unified account of the way things move, a beautiful theory of force and motion as it applies on earth and in space. His also did enormously important and influential work in optics, the study of the behaviour of light. Trinity College Chapel has a wonderful statue of Newton that is larger than life and towering overhead in a way that seems designed to encourage worship. In his hand, he holds a prism representing his research in optics.

Most scientists do not take much interest in the philosophy of science; they get on with their work without being particularly reflective or self-conscious about what they are doing. But some scientists do stand back from their own practice and attempt to understand better what they are doing, how the work should be done and what it can be taken to achieve. Newton was such a scientist. He made important contributions to the philosophy of science. Indeed it appears that Newton's philosophy actually influenced his scientific practice; certainly he used philosophical arguments against his scientific opponents. This rather goes against what I suggested earlier about the general lack of influence of the philosophy of science on science, but then Newton was an exceptional scientist.

One philosophical dispute in Newton's time concerned the question of whether we should understand scientific theories as revealing hidden truths about the world, the realist view, or instead take them to be more like computers, whose purpose is calculation rather than description, the instrumentalist view. The point of theories on the instrumentalist view is not to describe a hidden reality, but to provide tools for calculating accurate predictions of the observable world. The dispute between realists and instrumentalists is a philosophical perennial and remains a central topic in the philosophy of science today. Newton took a clear stand on the dispute. He held that science should be in the business of uncovering real causes, the truth behind the appearances, and he is important in the philosophy of science partially because of the way he promoted this realist position.

Newton is also important in the philosophy of science because of the emphasis he placed on observation and experiment. Like the technological power of science, this is something we now take for granted, but it was not at all obvious to all the scientists and philosophers in Newton's time. He also realised the problem for science created by this dependence. If the reason for believing scientific theories is observation and experiment, then it seems that these theories can never be proven to be correct. The results of observation and experiment are never certain and, as we saw in the case of the black ravens, no amount of positive evidence will prove that a general theory is true.

Newton looked for a middle ground in his philosophy of science between conclusive proof or demonstration and mere conjecture. Demonstration, which may philosophers and scientists thought science ought to provide, is proof from self-evident first principles. Nice work if you can get it, but Newton realised that the role of observation and experiment rules it out. At the same time, he did not want a science that consisted of wild conjectures of merely plausible hypotheses. He sought a middle position, where although science does not generate the sort of proof pure mathematicians can provide, it is much more than guesswork. Newton claimed that somehow the data, could be generalised to a theory that deserves high confidence even if it remained forever unprovable.

This helps to make sense of Newton's most famous philosophical slogan -- `*Hypotheses non fingo*' -- `I frame no hypotheses' or `I feign no hypotheses'. This seems a very odd thing for Newton to have said, since he spent his life framing a great number of wonderful hypotheses. Newton deployed that slogan in a particular context, defending himself against the charge that his theory of gravity was unacceptable because, while it used gravity to explain many other things, it did not properly explain gravity itself. Newton admitted that in that sense he did not attempt to explain gravity, but that such an explanation would be merely hypothetical and in any event not required in order to justify the physics he did provide. In more general terms, what Newton meant by his slogan, I think, was that Bacon was right, that science has to start with careful observation, and that the theory had to in some sense emerge as a warranted generalisation of the evidence. When Newton said that he did not frame hypotheses, what he meant was that he did not simply invent hypotheses that would account for the data, but rather found a path from the data to the hypotheses, even though that path could never be a path of proof. Whether there actually is any such general path from data to theory remains a central question in the philosophy of science.

To reach the third of my four Cambridge contributors, we now jump from around 1700 to around 1850, though we stay at Trinity. This contributor is William Whewell. He also has a nice statue in the College Chapel, as does Francis Bacon, though neither of them can hold a prism to Newton's overwhelming figure. Still, Whewell was an extraordinary polymath. He did seminal scientific work on the motion of the tides and he was at various times Professor of Mineralogy, Professor of Moral Philosophy, Master of Trinity and Vice-Chancellor of the University. Whewell's great range of interests included questions of scientific terminology and he is, in fact, credited with coining the very word `scientist'. What is striking is how late he did it, around 1840. Before then, people like Newton would not have been called `scientists': they would have been called `natural philosophers'. Whewell was called the `Professor of Moral Philosophy' not because he was particularly interested in ethics, but because his chair was in philosophy rather than in science.

One of Whewell's central interests was the history and philosophy of science, and one of the reasons that he is such an important figure in HPS is because he did both the `H' and the `P'. In the Cambridge HPS Department today, we work hard not to let the history of science and the philosophy of science become separate intellectual islands. It may seem surprising that effort should be required to avoid this, but sadly it is what happens in many HPS departments elsewhere: the philosophers only talk to the philosophers, the historians only to the historians. The reasons for this are complicated, but one factor is the difference between the techniques of historical and philosophical investigation. Such a separation is however a terrific waste of intellectual potential, and our success here at Cambridge in bringing the two areas into productive interaction is a source of pride. Whewell is a model here, in the way he appreciated the importance of bringing the history and the philosophy of science together.

In his philosophical work, Whewell went against Newton, insisting on virtually the opposite of *Hypotheses non fingo,* at least as I have interpreted that slogan. Whewell claimed that scientists should be business of framing hypotheses in the sense that Newton proscribed: scientists should search for hypotheses that would unify the diverse evidence. Good evidence for a hypothesis is not just numerous and accurate, but also shows great *variety*, and that is one reason there is according to Whewell no simple path from that evidence to the hypothesis. Here Whewell emphasised an ancient idea about scientific understanding. On the surface the world is a mess, terribly complicated, because many different factors are interacting and we only see a small part of what is going on. Underneath the surface, however, we can find the fundamental forces, which may not be visible but which will reveal the unity and simplicity that underlies the superficial complexity. This ancient idea of unity beneath diversity has been enormously influential, and Whewell developed it in a particularly fruitful way.

The fourth of my four figures brings us into the twentieth century and moves us from Trinity to King's College. He is John Maynard Keynes, one of the most important and influential economists of the century. What is not so well known is that his first book, the book that got him his fellowship at King's, was on probability. This work is of particular importance to the philosophy of science because of the way Keynes understood and interpreted the notion of probability.

The question of how we should understand claims about probability is a central philosophical topic. One view is that claims about probabilities are really claims about statistical patterns. Thus, to say that a coin has a probability of one-half on this view is just to say that, if you were to toss the coin many times, it would come up heads roughly half of the time. Keynes, however, argued that this is not the fundamental notion of probability, which is instead a relation between claims, between statements. More specifically, Keynes held that probability claims are claims about the support that evidence gives to a hypothesis where the evidence does not entail the hypothesis (Newton's worry), but where the evidence makes the hypothesis more or less probable. This sort of question about the relation between evidence and hypothesis lies right at the heart of the philosophy of science; hence the great philosophical interest of Keynes work on probability.

We will return to Keynes, but first I want to consider what sort of common philosophical thread one might find running through these four Cambridge figures. In some ways what is more important is how they differed, intellectually and culturally. Certainly they had important philosophical disagreements, one of which we will shortly consider; but there is also an important common theme, the theme of empiricism. All four held that the fundamental source of knowledge about the world is observation. The contrast here is with rationalism, the view that it is fundamentally through thought, not observation, that we come to know how the world operates -- René Descartes, for example, of *cogito ergo sum* fame, was a great seventeenth-century rationalist. The British, however, have tended to be empiricists, and our four figures run to form. They all emphasised that science cannot be done exclusively in the armchair; that scientists have to get out to do the experiments (though Newton did some of his best experimental work very near his armchair, in his rooms at Trinity). But they also realised the price scientists have to pay for taking the empirical route to knowledge.

Newton is explicit about the need to do science empirically, to rely on observation and experiment. In the *Optics*, he wrote that `...although the arguing from Experiments and Observations by Induction be no Demonstration of General Conclusions; yet it is the best way of arguing which the Nature of Things admits of...' (1704, 404). Scientists have to forego the certainty that proof provides because proofs are not to be had for claims about the way the world works. The only way to discover how it works is through observation. But this creates a difficult problem, at least from a philosophical point of view, the problem of gauging the uncertainty. Once the idea of proof is abandoned and replaced by the idea that the evidence supports or undermines a hypothesis, the relation is one of degree. It is not a question of `Yes' or `No', but of `More' or `Less'. This is one reason that Keynes's work on probability is so important: probability is a measure of the more and less. The difficult problem that philosophers have worked on is to understand the factors that determine this degree. What sort of evidence supports or undermines a theory, and what makes for more or less support? By considering some of the factors that philosophers have suggested, we can set the stage for the philosophical puzzle that I want to consider in the final part of this paper.

Prediction and Prejudice

The factors that seem to increase the support for a scientific theory can be roughly divided into features of the evidence and features of the hypothesis or theory. On the evidence side, more supporting evidence is better than less. That is pretty obvious, but how much evidence the scientist has is not the only factor that affects support. Variety in the data is also an evidential virtue. A scientist who just repeats the same experiment over and over eventually reaches a point of diminishing returns, whereas a theory supported by a variety of experiments inspires greater confidence. Having accurate and precise supporting data is another evidential virtue, as is having the results of controlled experiments, where the scientist can be confident of the absence of disturbing influences. The same applies to so-called `crucial' experiments, where the evidence simultaneously supports one theory while undermining some of its rivals, and to evidence that would be very improbable unless the theory was true.

One can construct a similar list of theoretical virtues. One is the prior plausibility of the theory: how natural it is and how well it fits with other claims the scientists already accept. Simplicity is another theoretical virtue: the simpler theory is often given a better chance of being correct. Other theoretical virtues include the plausibility of the auxiliary statements that have to be used to wring testable consequences out of the theory and the absence of plausible competing theories. These lists of evidential and theoretical virtues should make it clear both why the support a theory enjoys is a matter of degree and also why philosophers of science find it so challenging to account in detail for the impressive but not very reflective way scientists test and evaluate their theories.

This list of evidential and theoretical virtues is intended to be relatively uncontroversial, but I want now to focus on a disputed factor, a factor whose epistemic importance is the matter of much debate among philosophers of science. The dispute concerns the contrast between successful prediction and `accommodation'. In a case of successful prediction, the scientist first has her theory and then goes on to deduce a claim about the outcome of an experiment or observation that has not yet occurred. She then makes the observation or performs the experiment and finds the predicted result. In a case of accommodation, by contrast, the scientist has the data in question *before* he constructs his theory, and proceeds to construct a theory around the data, ensuring that the theory he builds fits the data, accommodating the theory to the data.

The existence of this distinction between prediction and accommodation is granted by both sides of the debate: the issue is over its significance. The dispute is this: are predictions worth anything more than accommodations? In other words, should scientists give a theory more credit for its successful predictions than for its accommodations? Many theories will have both sorts of support to their credit: the theory will accommodate some data and predict others. The question is whether the predictions give a stronger reason to believe a theory than the accommodations.

In discussion, I sometimes try to settle the issue democratically, by having a vote. Members of the audience have three choices. First, they can vote for the claim that predictions tend to provide more support than accommodations; second, they can vote for the claim that the difference prediction and accommodation makes no difference, it does not matter when the data are known; or third, they can abstain, if they have no clear intuitions on the matter. The results of such votes are fairly consistent and rather interesting. Most people do vote, and so presumably have a view on the issue. Of those who vote, most vote that prediction is better than accommodation, but a large minority choose the second option, that it makes no difference. So the issue is controversial, and not just among professional philosophers of science.

For evidence of this disagreement among the professionals, we need look no further that the Cambridge people whose philosophical contributions we have been celebrating. If you think that there is something special about prediction, that it does tend to provide stronger support, then you will find William Whewell on your side. He wrote that `It is a test of true theories, not only to account for but to predict, phenomena' (1847, aphorism 39). That is as clear a statement as you could like that prediction has some special value over what Whewell calls `accounting' and what I have called `accommodation'. On the other hand, if you think that the distinction between prediction and accommodation makes no difference, you are also in excellent company. In the view of John Maynard Keynes, `the peculiar virtue of prediction or predesignation is altogether imaginary. The number of instances examined and the analogy between them are the essential points, and the question as to whether a particular hypothesis happens to be propounded before or after their examination is quite irrelevant' (1921, 337).

Which then is the right answer? People who think predictions are worth more than accommodations often say that accommodation involve building a theory around the data, that this is *ad hoc*, and therefore provides little support for the theory. But this is not a good argument. What does `*ad hoc*' mean? It is Latin, so it sounds sophisticated, but all it literally means is `purpose-built'. In this sense accommodation obviously is *ad hoc*: the whole point is to build a theory to fit the data. To say that it is *ad hoc* in its literal meaning is just to repeat that it is accommodation: it is not to say or to show that the theory is poorly supported or otherwise deficient. So on this reading, to argue that accommodating theories are *ad hoc* therefore they are poorly supported is to argue that accommodating theories are accommodating theories, therefore they are poorly supported, which is a *non-sequitur*, to use another Latin expression. On the other hand, the expression `*ad hoc* theory' is often used in English, at least by philosophers, in a derogatory sense that implies that the theory is poorly supported or otherwise unattractive. On that reading, the argument becomes that accommodating theories are poorly supported, therefore they are poorly supported. This is the opposite of a *non-sequitur* but equally flawed: it begs the question, assuming what was to be shown (in Latin, a *petitio principii*). Either way, the *ad hoc* argument fails.

Leaving the Latin behind, what other arguments are commonly given for the claim that predictions are better than accommodations? One is the argument from testing, according to which predictions are worth more than accommodations because it is only in its predictions that a scientific theory is tested, and it is only passing a test that gives a scientific theory genuine credit. The idea is that a test is something that could be failed, and it is only a prediction that a theory can fail. In that case, the theory is made to stick out its neck in advance and say how things will be, so that the scientist may go on to discover that things actually are not that way. So, if the theory passes this test, and things are found to be the way the theory said they would be, the theory deserves some credit. In accommodation, by contrast, the theory does not stick its neck out, it cannot be shown to be wrong, because the theory is constructed after the data and compatibility is guaranteed in advance.

An analogy helps to bring out the intuitive strength of the argument from testing (cf. Nozick, 1983, 109). Suppose that Jacob, my elder son, takes his trusty bow and arrow, shoots at a target on the side of a barn, and hits the bulls-eye. We are impressed and give him a lot of credit. Now Jonah, my younger son, steps up to a different barn, pulls back his bow and shoots his arrow. Then he walks up to the side of the barn and paints a bulls-eye around his arrow. We would give him rather less credit, for archery anyway. That is the idea behind the argument from testing. Accommodation is like drawing the bulls-eye afterwards, whereas in prediction the target is there in advance. This argument seems clearly to show why successful prediction should count more than accommodation.

Nevertheless, as it stands this too is a bad argument. It confuses the scientific theory with the scientist, the theory with the theorist. What is true is that only in the case of prediction does the scientist run the risk of getting egg on the face; it is only in the case of prediction that the scientist may have to admit to having made a false prediction. But we care about the theory here, not the scientist, and from the point of view of the theory the contrast between prediction and accommodation disappears. If the predicted data had been different, that theory would have been refuted or disconfirmed, but just the same goes for accommodated data. If those accommodated data had been different, the theory that was built around it would also have been refuted. It is also true that, had the accommodated data had been different, the scientist would have built a different theory, but that is not to the point. From the point of view of the theory, the situation is exactly symmetrical. So the argument from testing fails.

Perhaps then Keynes was right and the supposed advantage of prediction over accommodation is imaginary. Many philosophers of science agree with him. Nevertheless, I will end this essay by suggesting where one might look for cogent arguments for the superiority of prediction. There are two promising types of argument. One is relatively straightforward; the other is a bit like trying to scratch the right ear with the left hand.

The relatively straightforward argument is the argument from choice. It depends on the fact that scientists can often choose their predictions in a way that they cannot choose which data to accommodate. When it comes to prediction, they can choose their shots, they can decide which predictions of the theory to check. Accommodated data, by contrast, is already there and scientists have to make what they can out of it. But how can this be used to show that predictions tend to provide stronger support that accommodations? A scientist will wish to make the strongest case to the scientific community that his theory is correct. So he has a motive for choosing predictions which, if correct, will give maximum support to his theory, not because they are predictions, but because they will exhibit the sort of evidential virtues mentioned before. Thus the scientist will choose predictions that allow for very precise observation, which would substantially increase the variety of data supporting the theory, and so on. Scientists will tend to choose predictions that will provide more support than in the case of accommodation, not directly because they are predictions, but indirectly because scientists have control over which predictions to check, control that is not available in the case of accommodation.

That is the straightforward argument from choice. It is probably cogent, but it does not show quite as much as one might hope. It shows why predictions *as a whole* tend to be more powerful than accommodations, but it does not give a reason for the more ambitious claim that a single, particular observation that was accommodated would have provided more support for the theory in question if it had been predicted instead. To try to make out that claim, we need a less straightforward argument, the fudging argument. It is related to both the *ad hoc* argument and the argument from testing we considered before, but it may avoid their weaknesses.

The fudging argument depends on an interesting feature of the lists of virtues, namely that some of the evidential virtues are in tension with some of the theoretical virtues. Here is an example. On the evidence side, scientists want the supporting evidence to be extensive and varied. On the theoretical side, they want the simplest theory. It is easy to have either one of these virtues on its own. If one just want lots of varied evidence one can just collect an encyclopedia full of facts; but the `theory' that is their conjunction will be incredibly ugly because the facts are so heterogeneous. On the other hand, if all that matters is simplicity, that too is easy, so long as one doesn't mind about fitting any of the evidence. What is hard and what scientists want is simultaneously to satisfy both constraints. They want a simple theories that nonetheless handle a great diversity of evidence.

Now for the fudging argument. When scientists have data to accommodate, they do the best they can. If the data are diverse, however, this can lead to a sacrifice in simplicity and other theoretical virtues. That is what I mean by `fudging': the scientist may, perhaps subconsciously, fudge the theory, putting in a few epicycles or extra loops to ensure that more of the data gets captured. In a case of prediction, by contrast, the scientist has no motive to introduce anything unnatural into the theory, because she does not know the right answer in advance and so would not know what kink to introduce into the theory even if one were required. So in this case the scientist will use the simplest theory and, if the prediction is successful, will have exercised both empirical and theoretical virtue.

The advantage that the fudging explanation attributes to prediction, is a bit like the advantage of a double-blind experiment that a doctor might perform to test the efficacy of a new drug. In a double-blind experiment, neither the doctor nor the patients knows which patients are getting the placebo and which are getting the drug. The doctor's ignorance makes his judgement more reliable, since he does not know what the `right' answer is supposed to be. The fudging argument makes an analogous suggestion about theoreticians. Not knowing the right answer in advance -- the situation in prediction but not in accommodation -- makes it less likely that the scientist will fudge the theory in a way that makes for a poor support. If you think about the puzzle of prediction and accommodation for yourself, as I hope you will, you may think of some objections to the fudging argument, but the argument may give one of the reasons predictions can be more valuable in science than accommodations, one reason why, on this issue, Keynes was wrong and Whewell was right.

Peter Lipton is Professor of the History and Philosophy of Science and a Fellow of King's College, Cambridge.

Further Reading and References

The philosophical work of the Cambridge figures mentioned in this essay are discussed in Fisch and Schaffer (1991), Gower (1997), Losee (1993) and Quinton (1980). Two good anthologies of recent work in the philosophy of science are Boyd, Gasper and Trout (1991) and Papineau (1996).

Richard Boyd, Philip Gasper and J.D. Trout, eds. (1991) *The Philosophy of Science*, Cambridge, MA: MIT Press.

Menachem Fisch & Simon Schaffer, eds. (1991) *William Whewell: A Composite Portrait*, Oxford: Oxford University Press.

Barry Gower (1997) *Scientific Method: An Historical and Philosophical Introduction*, London: Routledge.

John Maynard Keynes (1921) *A Treatise on Probability*, London: Macmillan, 1973.

John Losee (1993) *A Historical Introduction to the Philosophy of Science*, 3rd ed., Oxford: Oxford University Press.

Isaac Newton (1704) *Optics*, New York: Dover Publications, 1952.

Robert Nozick (1983) `Simplicity as Fall-out', in L. Cauman *et al.* (eds.) *How Many Questions?*, 105-119, Indianapolis: Hackett.

David Papineau, ed. (1996) *The Philosophy of Science*, Oxford: Oxford University Press.

Anthony Quinton (1980) *Francis Bacon*, Oxford: Oxford University Press.

William Whewell (1847) *The Philosophy of the Inductive Sciences*, 2nd ed., 2 vols., London: J.W. Parker.