

HOW CAN WE EVER KNOW WHAT WE KNOW?

Bryan Magee

The conception of knowledge that human beings have – what knowledge is, how we get it, and how we can be sure it is knowledge – is something that has always been changing. It changes historically, and it changes in the personal life of each one of us. It is changing even as I write, and will continue to change.

The conception of knowledge that human beings have – what knowledge *is*, how we get it, and how we can be sure it is knowledge – is something that has always been changing. It changes historically, and it changes in the personal life of each one of us. It is changing even as I write, and will continue to change.

Each of us individually goes through a stage of early development in which what we know, and what we think we know, are either what we have experienced directly or what we have accepted from others. Little actual thinking goes on. In infancy, what we learn from others comes usually first from our parents and other relations, and then from other children. When we start going to school we learn things from teachers of a quite different sort. When we start to read we learn new things from what we read. For quite a long time we tend to treat all these sources as authorities: we know because they have told us.

In societies also, most knowledge is handed down from one generation to the next, and societies, like individuals, go through periods of early development. In the more undeveloped ones it is assumed that the only way to be sure about anything is to get it from the proper authority. Outside the family this could be a social or religious

authority (the two could be the same), usually via a representative. It could also be a dominating group or class, or form of government, or even a weighty social consensus. The ultimate authorities could be dead – society may be bound by traditions carried over from the past. If the ultimate authorities are religious then the ultimate truths are known to be true because a priest or holy man says so, which means that a church or other religious institution says so, or a sacred text. If the sacred text is believed to have been dictated or inspired by God it may be socially impossible to dissent from it, perhaps even to question it.

In individuals first, then in societies, a new stage may develop at which some people start questioning authority. As individuals we come to realize that not everything our parents told us is true. This may be disconcerting, sometimes disillusioning. We also discover that some of the things our teachers told us are not true. Certainly not everything we have read is true. People often develop doubts about whether the religion they have been taught is true. In short, we realize that a thing is not necessarily true because it comes from an authority, however highly respected. So from then on, authorities will no longer be seen as infallible, even though they will continue to have a considerable importance in our lives.

Naturally, the authorities do not like being dissented from. Parents get angry when their children dispute what they say; and so, sometimes, do teachers. Organized religions take it very badly when their teachings are publicly rejected. In undeveloped societies it is not permitted for individuals to reject authority publicly; the penalties for doing so are severe, often lethal. A society needs to have reached an advanced stage before a dissenting individual can say publicly, in peace and freedom, what he really thinks, and teach his dissenting views to others, and publish them. Even in today's world it is still the case that in most societies – including the biggest of all, China – such things are not allowed.

The first society we know of in which those freedoms emerged was the world of ancient Greece. Teachers arose who, instead of giving their pupils a body of doctrine to be accepted unquestioningly by them and handed on in the same way to others, encouraged discussion and debate. The most remarkable thing about those teachers was that they did not require even their own pupils to agree with them. This was a revolutionary advance in human development. Instead of people accepting the 'truth' uncritically from an authority, the authorities themselves were teaching people to think for *themselves*.

In this way what we now call 'philosophy' began with the ancient Greeks. Some of those early Greek philosophers are still rated as highly as any – Socrates, Plato, Aristotle. Such figures introduced not only the principle of arriving at knowledge by the independent use of reason but also the principle of subjecting it to public criticism. Dialogue and debate were developed by them as methods of getting at the truth. They taught that if anyone came up with a good idea the way to test it was to think it through in the privacy of one's own mind and then subject it to public criticism. They taught their pupils to voice all the seemingly reasonable objections that could be brought against an idea and then try to think of answers to them. Socrates developed a teaching method based on searching interrogation which is still used, and still known as 'the Socratic method'. Aristotle tried to identify the differing forms that arguments can take, and to discover, by logical analysis, which of them are valid and which invalid. In this way he worked out a complete system of logic intended to cover all cases. For two thousand years, when people studied logic it meant Aristotle's logic. These were stupendous personal achievements.

Yet in those people's notions of how to arrive at new knowledge something vitally important was still missing. Let us pause to consider how things stood at that point. It had always been the case that most of our basic knowledge was handed down from the past, via parents, teachers, written sources, leaders, and other established authorities.

But we had learnt not to accept this uncritically, and to realize that some of it might be mistaken. We had learnt to evaluate it by the use of our reason, to think about it logically, question it, analyse it, argue with others about it, try to get at how things really are, and to do this not only in our private thoughts but in public discussion and debate. Some of the greatest advances that have ever been made were made by these methods, with their limitations, and went on being used successfully for something like two thousand years. But a crucial development still lay in the future.

Up to this point, everything had been in the realm of ideas, publicly expressible in language. The vital step that remained was to confront these ideas not only with dissenting ideas from other people but with the reality the ideas were supposed to explain. Plausible explanations, however attractive and persuasive, are not always correct. They need to be tested against the facts. We managed to come a long way confining ourselves to the examination of ideas alone. By intellectual analysis we were able to establish whether an argument or body of ideas is logically consistent – if it is not, it cannot be right. However, if it is consistent that does not mean that it *is* right, it means only that it *could* be right. It is right if, and only if, it corresponds to reality. That is something we can discover only by checking reality as well as ideas. When we investigate both we find that sometimes they correspond and sometimes they do not. Most intelligent and educated people nowadays are aware that ideas may be logically coherent, also attractive and persuasive, and yet still be mistaken. But it was a long time before this fundamental breakthrough in understanding was made.

Aristotle, who certainly had some sort of scientific bent, says somewhere in his writings that men have more teeth than women. This seemed to him to stand to reason. On the whole, men are physically larger than women, they have bigger heads, bigger faces, bigger jaws, so it would seem only natural for them to have more teeth. But they do not, they have the same number of teeth. Aristotle was

married twice, and it would have been easy for him to ask one of his wives to open her mouth while he counted her teeth. But obviously – and this is a prodigiously important point, especially since it is so difficult for us now to understand – it did not occur to him to do this. He just took it for granted that the right way to acquire knowledge was to think things through carefully, reason them out, if possible from first principles. But in this case (as in so many others) what seemed a natural assumption until it was checked against the facts turned out to be mistaken.

It seems puzzling now that it took Western man so long to realize the indispensable need for this last step. Yet perhaps, at that, it should not be all that surprising. For it is a common human failing to have unrealistic beliefs, attitudes and expectations, ideas that we are attached to but which, for one reason or another, do not correspond to reality. People can believe utterly mistaken things about their own society, surrounded even though they are by the reality of it every day, with the facts staring them in the face. I suspect most of us are familiar with this in some other people: we might even, when younger, have been like it ourselves. The truth is it requires self-discipline not to be.

It was not until the seventeenth and eighteenth centuries that this new way of thinking took hold. It happened first in Western Europe and brought about a sea change in Western thought. Among other things, it saw the beginnings of what we now call science. The new science began, surprisingly, with observation not of the things closest to hand but of those farthest away, the heavenly bodies. It brought to bear a new rigour in everything to do not only with observation but also with the drawing of logical consequences from observation. And it produced results with astonishing speed. At first these results were shocking, because they contradicted what had been believed for hundreds or thousands of years. For instance, in the early seventeenth century one of the founders of modern science, Galileo, published the assertion that the earth was revolving round the sun and at the same time

rotating on its axis. Actually these ideas, purely as ideas, were already nearly a hundred years old, because they had been offered tentatively once before as hypotheses. When, earlier, Copernicus had been attacked for being the first to publish them, he defended himself by protesting that he was not asserting any of this as reality, he was merely playing with the mathematics of it. It was all just speculative, he said. But now Galileo was asserting it as fact. The reason why this was so shocking was that it directly contradicted what was said in the Bible, where the earth is said to be fixed and immovable (Psalm 93) and the sun is described as going round the earth (Book of Joshua). Galileo was publicly condemned by the Pope. In conditions of serious personal danger he was forced by the Inquisition to recant. But later in the selfsame century, in 1687, Isaac Newton published in Protestant England an accurate working model of our whole planetary system.

One would have expected science to begin by applying this new kind of observation to objects close at hand. But one thing that is special about the stars as objects of observation is that human beings cannot interfere with them. We are, because we cannot help being, detached observers. We can note and measure and calculate, but we cannot – as yet, anyway – do anything to influence the stars, or change them, or move them around. We cannot *experiment* with them. We have to accept them as given to us, at least in the appearances we receive of them. But as soon as people applied the new ‘scientific’ methods to things nearer home it became a different story. If Nature did not present us, close at hand, with circumstances ideal for our observations and measurements, we might be able to set them up artificially. This led to what became known as ‘the experimental method’. Its beginnings, as one would expect, were simple, one might almost say simple-minded. For instance, in the seventeenth century Robert Hooke found himself wondering if an object would weigh the same at different altitudes. His reasoning led him to suppose that the further away an object was from the earth’s surface the

weaker the force of gravity on it would be, and therefore the less it would weigh. This happens to be correct, so reasoning alone had led him to the right answer. In previous ages he would have stopped there. Earlier thinkers such as Aristotle would have regarded him as having achieved a new piece of knowledge by good thinking. But Hooke knew that although reasoning had led him to this plausible conclusion, that was not enough to establish its validity. This could be done only by checking his reasoning against observable fact. So he thought up a way of doing this. He climbed to the top of Westminster Abbey carrying some scientific scales, together with a long, strong thread and a piece of iron. At the top he weighed the thread and the piece of iron, then let the iron hang down from the scales by the thread until it was just clear of the Abbey floor, then he weighed them again. The point of doing all this was that he knew his theory might be wrong. This had happened quite recently to a theory that had similarities with his. Until shortly before his time, everyone had believed that the heavier a physical object was, the faster it would fall. That seemed obvious: it was what common sense would have led anyone to suppose. It stood to reason. But Galileo had made the startling discovery that all bodies fall at the same velocity regardless of their weight, unless of course they are interfered with by other forces. *That*, you may say, is contrary to common sense, contrary to reason. But it is also true. And a lot of other basic truths about our world are contrary to common sense. For instance, the fact that we live out our whole existence on the surface of a giant ball which is hurtling through space at a speed of thousands of miles an hour and at the same time rotating on its axis. Nothing in common sense or even common observation suggests this. For most of human history it does not seem to have entered anyone's mind. The first people to suggest it were looked on by others as fantasists, saying something so crazy it could not possibly be true.

Many of the discoveries of science are extraordinary in this way, things that almost no one would have dreamt of,

at least not until shortly before they happened. If you ask yourself why everyone believed for thousands of years that the sun was going round the earth, and not that the earth was rotating, the answer may seem obvious. It *looks* as if the sun is going round the earth. You actually 'see' it doing so. There it is: it comes up in the east every morning, and with your own eyes you see it circle through the sky until it goes down in the west every evening. And it has to go round the earth underneath, because it comes up in the east again next morning. Nothing could be more self-evident. But if you now ask yourself: 'Well how would it have looked if it had looked as if the sun was not going round the earth but as if the earth was rotating?' you realize with a shock that it would have looked exactly the same. What is more, it is the second explanation that is the right one. Yet no human being seems to have thought of it for tens of thousands of years. To be the first to think of such a thing you need to be an extraordinary person, and to have an extraordinary mind. Making such discoveries requires not only high intelligence but also independence, the ability to think what no one has thought before. It calls for courage as well as creative imagination.

Three or four such individuals in the seventeenth century played decisive roles in getting educated people in the West to understand and embrace the possibility of what we now call scientific knowledge. Galileo and Newton were two, another was Descartes. He was the best philosopher among them, and also an original mathematician, the inventor of analytic geometry and the graph. He was especially impressed by what seemed to him the unique certainty of mathematical knowledge. Other people were too, but what Descartes thought of doing was analysing the methods of mathematics to discover what gave it its certainty, and then applying those methods in other fields. This is why his masterpiece is called *Discourse On Method*. He came to the conclusion that what gave mathematics its unique reliability was as follows. A mathematical proof starts from the minimal number of premises, and

these premises are of the utmost simplicity. They are so basic, so obvious, it seems impossible to doubt them – for instance, that a straight line is the shortest distance between two points. From premises such as these a chain of mathematical reasoning moves slowly forward by elementary logic, one small step at a time, each step not only transparent but impossible to refute, until we find ourselves reaching conclusions so distant from our starting point that they had not been in the least obvious when we began. Worlds of unexpected knowledge are opening up in front of our eyes, some of it of immense practical usefulness. And it had all been reached by impeccable logic from premises it was impossible to doubt.

This, said Descartes, is the right way for us to acquire empirical knowledge of the world around us: start from facts of which we can be really sure, then everything, *everything*, that we deduce from these facts by strict logic must be true, however unobvious it may be. He was putting forward a programme for acquiring what would now be called scientific knowledge, but the concept of science had not yet come into existence. Unlike Galileo and Newton, Descartes did not give us any important scientific knowledge, but he told us how to get it. And for hundreds of years his method (or something very like it) was accepted as the basis of scientific procedure. It was spectacularly successful. This new science began to transform the world we live in. But each half of Descartes' method, taken separately – factual observation and logical deduction – was a source of problems that led eventually to his view of scientific method being superseded by another. Just like the wonderful advances of the Ancient Greeks, his approach was eventually to be replaced by a better one.

The most important practical limitation of logic is that it cannot, by itself, prove the truth of anything. Sophisticated thinkers have long understood this, but I have the impression that the generality of people still do not. They think that logic proves the truth of things. But even the most logically impeccable proof does not prove the truth of its

conclusion. It proves only that its conclusion follows from its premises. It could be that not all its premises are true. How can we be sure they are? The truth of the premises cannot be proved by the argument itself because they need to be presupposed before the argument can begin. An argument that tried to prove its own premises would be assuming what it set out to prove. That would be circular, and a vicious circle. You may say that its premises have been established as the conclusions of earlier arguments, rigorously conducted, but those earlier arguments must have had premises whose truth they did not prove. We are in an infinite regress here. Every rational argument, every proof, every demonstration whatsoever, has to have at least one premise before it can begin, and the truth of that premise cannot be proved by logical argument, however far back we pursue it. What we need at the very beginning is at least one premise which is not the conclusion of a logical argument, in other words a premise known to be true *regardless* of logical argument. Only if that premise is true can all, or any, of the conclusions logically derived from it be wholly relied on. It is also, of course, necessary for the logic itself to be impeccable, but that is never a sufficient condition. If the logic is good the burden of sufficiency is shifted to the reliability of the premises. But if that cannot be established by logic, how can it be relied on?

This is where, historically, direct observation came in, as the provider of indubitable facts which could then function as initial premises for chains of deductive logic. Observation would need to be conducted by certain rules if it was to be dependable as a method, but most of these rules were obvious, and were soon absorbed into both practice and theory. Direct observation must take place in carefully controlled conditions in which the possible involvement of variables is clearly understood, and all the relevant variables controlled – for instance either removed or standardized. Everything that can be measured must be measured with care, then double-checked. Every procedure must be repeated more than once to make sure it always

yields the same result. Then it must be gone over again by other people to make sure the first person was not persistently making the same mistake, or using faulty equipment, or under the influence of subjective factors of which he was unaware. Although no human observation, or activity of any kind, can be relied on to be one hundred per cent objective, this came to be provided for in the practice of science, and complete objectivity was replaced as a goal by intersubjectivity, mutual checking, that could be made as exhaustive as any human activity can be. This had the consequence of making science a social activity within which the exchange of undistorted information was indispensable. Thus the development of rationality and science had profound implications for developments in society, and these have been important ever since.

After Newton, one of the chief aims of innovative scientific activity became the discovery of so-called scientific laws, laws of nature. These are not descriptions of unique occurrences but general statements that tell us universal truths about how things are – for instance the inverse-square law, which tells us that every material object attracts every other with a force that is directly proportional to its mass and inversely proportional to the square of the distance between the two objects. This gives us an amazing amount of detailed knowledge about the universe in which we find ourselves. Although such scientific laws are unrestrictedly general they have often, in practice, been arrived at because individual observations pointed us in their direction. Over and over again when situations of a certain kind were observed and measured, and changes in them observed and measured, they were found to display such-and-such a characteristic - until people began to expect that this kind of situation was always going to have that characteristic. A scientist would then frame a general hypothesis to this effect and try to think up a crucial experiment to, as he would put it, prove the truth of his hypothesis - thus making him the discoverer of a scientific law. For a long time this was regarded as foundational to good scientific method.

Laws of nature were thought of as unrestricted generalizations arrived at from direct observations by a process called inductive logic. However, there was a flaw in this that was to prove fatal. The flaw had long been evident to some of the wiser heads, not only among scientists and philosophers but even people in everyday life.

For we *do* use this method in everyday life. If one day I have a stomach ache after eating gooseberries I may not immediately associate the stomach ache with the gooseberries, but if it keeps happening – if every time I eat gooseberries I get a stomach ache – I make the inference that eating gooseberries gives me a stomach ache. Huge amounts of our small-scale, everyday knowledge about ourselves and other people and the world around us are arrived at in this way, by forming generalizations based on repeated experience of the same individual instance. But the inconvenient truth is that it is never reliable logically to base unrestricted generalizations on single instances, no matter how numerous the instances. From the fact that every swan any known person has seen is white it does not follow that all swans are white: in some as-yet unexplored part of the world, swans may be of another colour. I use this example because it is something that actually happened. For thousands of years every swan any Western man had ever seen was white. In the late Middle Ages a commonly used textbook gave ‘All swans are white’ as an example of a known truth. Yet after all those thousands of years, and all those millions of observations, crowned by the authority of a textbook, when Europeans discovered Australia they discovered black swans. It is impossible to exaggerate the importance for human knowledge of this example. It shows that from no number, however large, of true observations can we infallibly infer an unrestricted generalization. And the significance of this for science is seismic, because most scientific laws are unrestricted generalizations supported by a finite number of observed instances, and it simply is a fact that no matter how many of these individual instances are confirmed, no general

truth *logically* follows from them. We have, it seems, a natural tendency to generalize from individual experiences, but if this is so it is a psychological process, not a logical one. We are making an association of ideas, not a logical connection. There is no such thing as inductive *logic*. Some of the great early-modern philosophers perceived this, for instance Locke and Hume, but it was not until the twentieth century that people more generally woke up to it - and to what it meant for science. It meant that scientific laws had not been proved. What is more, they never *could* be proved. They could not be proved by logical deduction from past observations, and they could not be proved by direct observation either, because it would always be impossible to observe future events. The most that could be said for these so-called laws is that they had withstood all the practical and theoretical tests to which they had ever been subjected. But they still had to be regarded as fallible, and therefore provisional: after all, the assertion 'all swans are white' withstood all tests successfully for thousands of years.

This opened people's eyes to why it is that scientific revolutions keep on occurring - indeed, to how it is that scientific revolutions can occur at all. None of our scientific knowledge is conclusive, definitive. It can always be changed for the better, always improved on. What we are doing at any given time is using the most reliable knowledge we have; and we go on using it until it runs up against either a direct observation that contradicts it or a problem we cannot solve with it. Such a problem indicates to us where there is a mistake in our present understanding, and we either revise our theory or abandon it altogether. In either case we exchange it for a better theory. In my lifetime alone all the major sciences have gone through revolutionary conceptual upheavals. And this will recur indefinitely into the future.

It is essential to realize that the fact that we cannot prove the truth of a scientific theory does not mean that one theory is as possible as another, still less that we are free

to believe what we like. Although a theory cannot be conclusively proved it can be conclusively disproved. Although the statement 'all swans are white' can never be conclusively proved by any number of observations of white swans, one single observation of a black swan disproves it. So although a general empirical theory cannot be verified it can be falsified. And this means it can be checked – tested against reality, and possibly found wanting. The example of swans is an especially good one because, if Australia had never been discovered, Western man could have gone on for ever regarding it as certain that all swans were white. He would never have had any reason to think otherwise. Even as things turned out, he went on believing it for thousands of years. In the same way in all sciences we make fruitful use of 'the best of our knowledge' until we run up against contradictions, and then we have to think again.

Perhaps the biggest field of familiar examples to us today is medical science. It seems as if scarcely a day goes by without some advance in it showing that procedures we have been using up to now are not the best ones, and may rest on a false assumption. And we are familiar with the fact that the new methods will themselves be overtaken in the course of time – and after that be overtaken again. We also understand that although final certainty will never be reached, this does not mean that any theory is just as good as any other, or that we are free to adopt any methods we like. A hospital that took that approach would kill more patients than it cured.

So relativism is not an option. Despite the fact that we are forced into a fallibilist view of knowledge, an attitude of 'my 'truth' is just as likely to be valid as your 'truth' is completely illegitimate. For we can and do have grounds for regarding one theory as better than another: while the *truth* of a general theory cannot be justified, a *preference* for one theory over another can be. And that is the death-knell for relativism. It is also the death-knell for an attitude that confronts all theories with an equal scepticism, for the rational

approach I have outlined does as a matter of brute fact lead to improvements in knowledge, and therefore it supports an optimistic attitude towards the growth of knowledge.

In physics there was not one major revolution in the twentieth century but two: relativity theory and quantum theory. Historically it was these more than anything else that brought home to people that scientific knowledge was not a body of unchanging fact. For two or three hundred years, they had supposed that it was. Classical physics, Newtonian physics, seemed to provide us with definitive knowledge. Newton's so-called Laws were seen as, literally, Laws of Nature, and were taught under that name: they were accepted as factually true descriptions of how Nature worked. In the words of Alexander Pope:

Nature and Nature's laws lay hid in night:
God said, Let Newton be! and all was light.

However, in the twentieth century Einstein came along and showed that Newton's physics, though brilliant, and phenomenally useful, was not accurate in every respect. It was possible to replace some of Newton's theories with better theories – and Einstein did. This required us to change our understanding of what knowledge itself was. It is not fact. It is not objective truth about the way things are. It is not unchanging and definitive. It is a useful approximation which can always be improved on; and because it is only an approximation it almost certainly *will* be improved on.

At first many people were puzzled and disconcerted by this. If our knowledge does not consist of facts, what does it consist of? The answer is, it consists of theories. But these are theories that have withstood the most rigorous testing, and done so better than any known alternatives. They are here for us to use until we find better ones. So the advance of scientific knowledge, which is taking place all the time, does not consist in the addition of new certainties to an already existing (and ever growing) body of certainties; it consists in the replacement of currently

employed theories by better theories. That is why we no longer refer to the historic breakthroughs in knowledge as certainties but as theories: even after a hundred years we are still talking of Einstein's theory of relativity, and quantum theory - and at last, accurately, of Newton's theories.

There is another way, too, in which examples of the fallibility of knowledge are offered by the two great revolutions in physics of the twentieth century. When the logical consequences of each of them are followed through it emerges that some of the logical implications of relativity are incompatible with some of the logical implications of quantum theory. This means that they cannot both be entirely correct. At least one of them is mistaken, and more probably both are. Not long ago there seemed to be evidence that Einstein's doctrine that nothing in the universe could travel faster than light was mistaken. At first, as usual, many people were disconcerted, and the evidence was disputed. But for as long as the matter remains unsettled we shall go on making use of the theory.

It is during the course of my lifetime that the view of knowledge that occupies most of the field today gained possession of it. The most influential scientist in this development was Einstein, the most influential philosopher Karl Popper. But a large number of other people played necessary roles. And the overall situation will continue, as it always has, to change.

If I were to sum up at this moment the best of our knowledge about how to get knowledge, it would go something like this:

We work through four main stages in order.

Stage one: a question is raised. It may be a practical problem, or we may just be curious, but for whatever reason there is something we want to know. We can call this our problem, or our problem-situation. We would be well advised not to hurry ahead from this too quickly but to

turn it over and think it through, because usually our success in solving a problem will depend partly on the accuracy, clarity and depth with which we have understood it – and understood not only it but also its implications. So we should acquire a thorough, all-round grasp of the problem before attempting to move forward from it.

The second stage is to move forward, and search for a possible solution to the problem, a solution that will genuinely work if it is a practical problem or, if it is theoretical, one that meets all the objections that we and others can throw at it. Creative imagination plays a role here, and so does independence of mind. Most of the great advances in knowledge have a boldness and freedom that few people are capable of, so new ideas tend to get a mixed reception from other experts – which is good, because the new ideas are then scrutinized by specialists looking for faults. At this second stage we have a proposed solution whose preferability is only a hypothesis: not everyone accepts it. It cannot be judged by how it was arrived at, because it may have been no more than a hunch. Whatever its origins, a hypothesis will be judged not by how it was arrived at but by how it responds to tests.

This carries us on to the third stage, the testing of the proposed solution. This involves searching for – and then, if necessary, setting up – conditions in which the proposed solution can be set directly against observed reality, which may show it to be false. This is the role of experiment. It, too, can call for ingenuity. Testing is not necessarily a matter of black and white, yes or no. Experiments can bring out strong as well as weak points in a theory. In the course of showing parts of our approach to be wrong they may at the same time indicate ways in which these can be improved. So our proposed solution may have some of its errors eliminated by the experimental process, and be strengthened thereby.

By the end of all this we are in our fourth and final stage. We now either have a solution to our initial problem or we do not. If we have, the initial problem-situation no longer

exists, and having solved the problem we are in a new situation. But the new situation raises new questions which had not presented themselves before. So we have not reached a stopping place, we just have new and different problems. In this manner we find ourselves moving forward in an endlessly ongoing process. The growth of our understanding and of our knowledge never ends, but at the same time our ignorance grows with our knowledge. Every advance gives rise to new questions. Our understanding is like the circle of light cast by a lamp in the darkness: when the circle grows, its frontier with the darkness lengthens.

This is what happens when an attempt to increase our knowledge is successful. But most of our attempts are not successful. Most of our good ideas do not survive the tests to which we subject them. Before we embark on putting them into effect, common sense and worldly wisdom have already warned us to expect snags, unforeseen consequences, unwelcome side-effects. But even outright failure teaches us something, and therefore adds to our knowledge. It eliminates what had seemed a promising line of enquiry. It teaches us that things are not as we had thought they might be; and it educates us further about this particular problem, deepening our understanding of it, especially where its greatest difficulties lie, and what the minimum conditions are that any viable solution will have to meet. Crucially, we learn from our mistakes, our failures, our disappointments. People who are unusually good at doing this are among the most creative and successful people there are. And for all of us, most of what we rightly call our experience, and value as such, consists of what we have learnt from our mistakes.

*

The developments I have outlined have led us to the best-attested knowledge we now have. This includes scientific knowledge, but not all knowledge is scientific knowledge: there are other kinds too. However, it has to be said that

the sciences provide us with our most publicly reliable knowledge, and also with a great deal of our most practically useful knowledge: medical science, engineering, the whole of modern technology – and, through those, modern industry and modern agriculture (one is tempted to say the modern world). Therefore a general theory of what knowledge is has to accommodate science if it is to be credible: if it does not apply to our best-attested knowledge it cannot be valid as a theory of knowledge. This is one reason why it is no longer possible for us to go on thinking of knowledge as justified true belief. Our scientific ‘knowledge’ cannot be justified; and if we cannot reach permanently acceptable certainty in our most reliable form of knowledge, how do we expect to reach it in any of the others? This realization, universal in its implications, is of historic importance.

Science is not alone in offering us knowledge and understanding in the form of explanatory frameworks that may approximate more or less to the truth. History does, as do the other so-called social sciences: sociology, economics, anthropology and the rest. So, indeed, does common sense. And so does metaphysics.

Metaphysics and science have a lot in common. Historically, it was out of metaphysical theories that most of science developed. The research programmes within which empirical investigations are conducted are usually metaphysical hypotheses. The chief criterion of demarcation between metaphysics and science is empirical testability, and we have said something about this. The concept of testability is a mixed one, partly logical and partly historical. A theory that is not at first empirically testable may become so through advances elsewhere, particularly in technology – and then a metaphysical theory becomes a scientific one. When it does, it is more often eliminated than not.

In metaphysics, as in science, human beings create hypotheses in order to make sense of the world and provide them with help and guidance. In both cases the hypotheses are formed mostly in response to either

practical problems or curiosity, and are intended both to explain the facts of our experience and to guide our expectations. Metaphysics and science are both constantly subjected to critical examination that reveals inconsistencies and self-contradictions which eventually lead to rejection of a hypothesis. In both, theoretical examination can by itself carry us a long way in assessing the relative merits of competing theories, and also the worthwhileness of giving time to their consideration. In neither, though, can we ever be conclusively sure that a theory is true, *even if it is true*, because we have no way of conclusively verifying it. For us, even the truth itself remains permanently open to question.

This abandonment of certainty as a prerequisite of knowledge, even as an attainable formulation for knowledge, has permanently altered the status of belief. When people were engaged, as for so long they were, in the pursuit of certainty, beliefs had a key role to play. They functioned as tentative or provisional certainties. In almost all areas of activity people would formulate beliefs (whether on the basis of experience, or insight, or hunch, or a combination of these) and then try them out to see if they were true. This happened in politics and all branches of practical life, including business. Researchers of every kind, from scientists to scholars, worked in this way more often than not. Even the task of philosophy was defined by major philosophers as the justification of our most important beliefs. But when certainty comes to be seen as unattainable, and the pursuit of it is abandoned, belief as tentative certainty no longer has a role to play. If there is no certainty there can be no provisional certainty. When 'not-knowing' is acknowledged as the only rational possibility, then so also is 'not-believing'. So belief is now superfluous, irrelevant. Worse than that, it gets in the way, it prevents us from understanding the reality of our situation. Whether we realize it or not, our so-called beliefs are conjectures. If well formulated, they may be true; but the fact that they *may* be true is no ground for believing that they are.

In everyday speech the word 'believe' is often used with the right degree of sophistication, like the word 'knowledge' in the phrase 'to the best of my knowledge'. If you ask me: 'Is Peter on holiday?' and I reply 'I believe he is,' you take this to mean not that I am sure of it but that I am unsure of it – that I have reason to think he is, and am proceeding on that assumption, but remain open to the possibility that he may not be, and am taking care to intimate this to you. This familiar use of 'believe' has the heart of the matter in it.

Our knowledge is our explanations, and in any pursuit of truth the best explanations we can have are well-supported hypotheses that are the least unlikely among the known alternatives. They may work out well in practice, as a basis for our actions. But far from actively believing them to be true, we need to be clear while using them that they may turn out to be wrong. Psychologically, this is incompatible with actively believing them to be true.

With the replacement of proof by progress comes the replacement of belief by conjecture. Belief, having in our new situation no path towards certainty, is an interim goal where there is no function for interim goals. Conjectures that are as imaginative and inspired, as well-informed and hard-worked-on as we can achieve, getting as close to the truth as we are at the moment able to come – Yes. In the formation of these the full deployment not only of information, experience and criticism but also of insight, imagination, creative thinking, speculation, guesses and hunches – Yes. But to be committed to a *belief* that a conjecture is true is misplaced, and is likely to misdirect our efforts, because it will weaken the critical attitude we need to bring to the conjecture and will narrow our openness to alternative possibilities. Thus it reduces our creativity, and also the level of our performance, and in both ways lowers our effectiveness in trying to get closer to the truth. To some of the creative uses we make of our conjectures we need to bring the negative capability that Keats attributed to Shakespeare: 'Negative Capability, that is, when a man is

capable of being in uncertainties, mysteries, doubts, without any irritable reaching after fact and reason.'

In the difference between the two approaches there are personal-character implications. When certainty is seen as attainable, and belief as a stepping stone towards it, persistent refusal to believe anything is self-defeating, because it prevents achievement of the goal. In such circumstances, the permanent withholding of commitment is a character defect, a lack of courage, an unwillingness to be decisive and move forward. But when certainty is exposed for the delusion it is, a permanently agnostic openness to alternative possibilities is the only legitimate approach. It is not an approach that treats all possibilities alike, or as having the same importance, but one that sees all as fallible. And it applies as much to metaphysics as to science. In the serious pursuit of understanding in any field, a determined clinging to positive belief is an opting out of the main task.

A well worked out theory of knowledge that rejected the search for ultimate foundations without relapsing into relativism or scepticism, indeed while offering rich explanations for the growth of knowledge and the success of science, appeared eventually in the work of Karl Popper. It was he who explicitly replaced the search for proof with the search for progress, first in science and then across all fields of knowledge. Instead of the metaphor 'foundations of knowledge' he suggested we use the metaphor of a house built not on foundations but on piles, as in some parts of the world they are. The piles have to be driven down deep enough to carry the weight of the structure, and if the house goes on being added to, the piles need to be driven deeper and deeper; but there is no line at which a limit to this process can be drawn. There is no ultimate level that will sustain the weight of any structure whatsoever. There are no ultimate foundations.

Everything I have written here drives home the symbiotic relationship between knowledge and ideas. But the ability to have good ideas is a rare one. I doubt whether many people have original ideas – most spend their time making

use (it may be good use) of other people's. At the opposite end of the spectrum there are a handful of individuals who are fountains of ideas: I have instanced Galileo, Newton and Einstein (and in other fields I could mention Shakespeare, Mozart and Michelangelo). But the ideas of these extraordinary people have their limitations too, and are sometimes mistaken, so they too can be criticized, and are capable of improvement.

Because all reasoning has to have ultimate premises which have not been arrived at by reasoning, there can never be a totally rational way of arriving at new ideas (as, for so long, people thought there was, by 'inductive logic'). Creativity cannot be wholly a rational activity. Good ideas are arrived at in a multitude of ways: perhaps most often by making unexpected changes in already-existing ideas, following up hunches or conjectures or guesses; but also in flashes of inspiration, or as a result of dreams or dreamlike states; even sometimes through misunderstandings and mistakes. At the deepest level one sees the rapt, sustained absorption of a creative person in the object of his concern, an interrelationship fusing receptivity with activity through his entire personality, not only his mind but his senses and emotions too, as if he were a creative artist. This has been superbly expressed by the most creatively original scientist of my lifetime, Einstein, who attempted in a now-published letter to Karl Popper to describe what he called his 'search for those highly universal laws ... from which a picture of the world can be obtained by pure deduction. There is no logical path,' he wrote, 'leading to these ... laws. They can be reached only by intuition, based on something like a feeling of oneself *into* the objects of experience.' In the same letter he expresses his agreement with Popper 'that theory cannot be fabricated out of the results of observation, but that it can only be invented.' He is here repudiating the belief that innovative theories can be arrived at by a so-called logical process called induction, using observation-statements as premises. (In insisting that there is no scientific or logical way of arriving at new ideas he is also

pointing to something fundamentally in common between the sciences and the arts.)

Perhaps a final word needs to be said about our use of the word itself, 'knowledge', which until recently was current in philosophy to mean 'justified true belief'. The point made by that usage was that, before I can claim correctly that I know something, three conditions have to be met. First, what I claim as knowledge has to be true: if it is not true it cannot be knowledge. Second, I have to believe it to be true: if I do not believe it, it cannot be part of my knowledge. Third, I have to have adequate grounds for believing it to be true: the belief must be not only true but justified. For instance, I might think I know that it is now three o'clock because my watch says it is, and my watch has always kept good time. But if, unknown to me, my watch stopped at three o'clock this morning, and by coincidence I look at it when the time is three o'clock in the afternoon, I am merely deceived into thinking I *know* that it is three o'clock. Pure chance is at work, deceiving me into thinking I know something when my belief is correct by the merest happenstance. It is not genuine knowledge - even though it is true and I believe it to be true. To be genuine knowledge my belief needs to have genuine justification.

When Socrates said, as he did repeatedly, that he did not know anything, it was this conception of knowledge that he was taking for granted. There have been other major figures in philosophy who similarly disclaimed knowledge. Locke said we have very little of it, and deal for the most part in probabilities. Hume thought we possess no certain knowledge outside the technical fields of logic and mathematics. A phrase I heard from Popper's lips many times was 'We don't know anything.' This realization has never been absent from Western philosophy. Even before Socrates, Xenophanes had it. But it was never a familiar view until twentieth century physics gave it scientific support. Until then, the prevailing view was that we may not have all that much knowledge as yet, for certain, but it is waiting to be found, and our task is to seek it out and get

as much as we can. But all this time most philosophers were using the word 'knowledge' not only in the absence of what they meant by it but in the absence of its possibility. And this had all the time been tacitly acknowledged in an everyday use of the term with which we are all familiar, the idea that we should act on 'the best of our knowledge' while at the same time regarding it as open to improvement. Wise administrators have always done this, as have successful men of business and great military commanders. The word 'know' already contains within itself these subliminal assumptions a good deal of the time, especially when we are going about our ordinary lives and doing the world's business. This is because that is the reality of the situation. Already in the philosophy of science, if not yet quite wholly in general philosophy, the battle for a conjectural view of knowledge has been won. It is not likely to be long before the fruits of this victory are generally shared.

Bryan Magee has reached audiences in philosophy via tv and radio as well as in print. He is an honorary fellow of Keble College, Oxford.