

Chapter 2

Scientific reasoning

Scientists often tell us things about the world that we would not otherwise have believed. For example, biologists tell us that we are closely related to chimpanzees, geologists tell us that Africa and South America used to be joined together, and cosmologists tell us that the universe is expanding. But how did scientists reach these unlikely-sounding conclusions? After all, no one has ever seen one species evolve from another, or a single continent split into two, or the universe getting bigger. The answer, of course, is that scientists arrived at these beliefs by a process of reasoning or inference. But it would be nice to know more about this process. What exactly is the nature of scientific reasoning? And how much confidence should we place in the inferences scientists make? These are the topics of this chapter.

Deduction and induction

Logicians make an important distinction between deductive and inductive patterns of reasoning. An example of a piece of deductive reasoning, or a deductive inference, is the following:

All Frenchmen like red wine
Pierre is a Frenchman

Therefore, Pierre likes red wine

The first two statements are called the premisses of the inference, while the third statement is called the conclusion. This is a deductive inference because it has the following property: if the premisses are true, then the conclusion must be true too. In other words, if it's true that all Frenchman like red wine, and if it's true that Pierre is a Frenchman, it follows that Pierre does indeed like red wine. This is sometimes expressed by saying that the premisses of the inference entail the conclusion. Of course, the premisses of this inference are almost certainly not true – there are bound to be Frenchmen who do not

like red wine. But that is not the point. What makes the inference deductive is the existence of an appropriate relation between premisses and conclusion, namely that if the premisses are true, the conclusion must be true too. Whether the premisses are actually true is a different matter, which doesn't affect the status of the inference as deductive.

Not all inferences are deductive. Consider the following example:

The first five eggs in the box were rotten
All the eggs have the same best-before date stamped on them

Therefore, the sixth egg will be rotten too


This looks like a perfectly sensible piece of reasoning. But nonetheless it is not deductive, for the premisses do not entail the conclusion. Even if the first five eggs were indeed rotten, and even if all the eggs do have the same best-before date stamped on them, this does not guarantee that the sixth egg will be rotten too. It is quite conceivable that the sixth egg will be perfectly good. In other words, it is logically possible for the premisses of this inference to be true and yet the conclusion false, so the inference is not deductive. Instead it is known as an inductive inference. In inductive inference, or inductive reasoning, we move from premisses about objects we have examined to conclusions about objects we haven't examined – in this example, eggs.

Deductive reasoning is a much safer activity than inductive reasoning. When we reason deductively, we can be certain that if we start with true premisses, we will end up with a true conclusion. But the same does not hold for inductive reasoning. On the contrary, inductive reasoning is quite capable of taking us from true premisses to a false conclusion. Despite this defect, we seem to rely on inductive reasoning throughout our lives, often without even thinking about it. For example, when you turn on your computer in the morning, you are confident it will not explode in your face. Why? Because you turn on your computer every morning, and it has never exploded in your face up to now. But the inference from 'up until now, my computer has not exploded when I turned it on' to 'my computer will not explode when I turn it on this time' is inductive, not deductive. The premiss of this inference does not entail the conclusion. It is logically possible that your computer will explode this time,

even though it has never done so previously.

Other examples of inductive reasoning in everyday life can readily be found. When you turn the steering wheel of your car anticlockwise, you assume the car will go to the left not the right. Whenever you drive in traffic, you effectively stake your life on this assumption. But what makes you so sure that it's true? If someone asked you to justify your conviction, what would you say? Unless you are a mechanic, you would probably reply: 'every time I've turned the steering wheel anticlockwise in the past, the car has gone to the left. Therefore, the same will happen when I turn the steering wheel anticlockwise this time.' Again, this is an inductive inference, not a deductive one. Reasoning inductively seems to be an indispensable part of everyday life.

Do scientists use inductive reasoning too? The answer seems to be yes. Consider the genetic disease known as Down's syndrome (DS for short). Geneticists tell us that DS sufferers have an additional chromosome – they have 47 instead of the normal 46 ([Figure 5](#)). How do they know this? The answer, of course, is that they examined a large number of DS sufferers and found that each had an additional chromosome. They then reasoned inductively to the conclusion that all DS sufferers, including ones they hadn't examined, have an additional chromosome. It is easy to see that this inference is inductive. The fact that the DS sufferers in the sample studied had 47 chromosomes doesn't prove that all DS sufferers do. It is possible, though unlikely, that the sample was an unrepresentative one.



5. A representation of the complete set of chromosomes – or karyotype – of a person with Down’s syndrome. There are three copies of chromosome 21, as opposed to the two copies most people have, giving 47 chromosomes in total.

This example is by no means an isolated one. In effect, scientists use inductive reasoning whenever they move from limited data to a more general conclusion, which they do all the time. Consider, for example, Newton’s principle of universal gravitation, encountered in the last chapter, which says that every body in the universe exerts a gravitational attraction on every other body. Now obviously, Newton did not arrive at this principle by examining every single body in the whole universe – he couldn’t possibly have. Rather, he saw that the principle held true for the planets and the sun, and for objects of various sorts moving near the earth’s surface. From this data, he inferred that the principle held true for all bodies. Again, this inference was obviously an inductive one: the fact that Newton’s principle holds true for some bodies doesn’t guarantee that it holds true for all bodies.

The central role of induction in science is sometimes obscured by the way we talk. For example, you might read a newspaper report that says that scientists have found ‘experimental proof’ that genetically modified maize is safe for humans. What this means is that the scientists have tested the maize on a large number of humans, and none of them have come to any harm. But strictly speaking this doesn’t *prove* that the maize is safe, in the sense in which mathematicians can prove Pythagoras’ theorem, say. For the inference from ‘the maize didn’t harm any of the people on whom it was tested’ to ‘the maize will not harm anyone’ is inductive, not deductive. The newspaper report should really have said that scientists have found extremely good *evidence* that the maize is safe for humans. The word ‘proof’ should strictly only be used when we are dealing with deductive inferences. In this strict sense of the word, scientific hypotheses can rarely, if ever, be proved true by the data.

Most philosophers think it’s obvious that science relies heavily on inductive reasoning, indeed so obvious that it hardly needs arguing for. But, remarkably, this was denied by the philosopher Karl Popper, who we met in the last chapter. Popper claimed that scientists only need to use deductive inferences. This would be nice if it were true, for deductive inferences are much safer than inductive ones, as we have seen.

Popper’s basic argument was this. Although it is not possible to prove that a scientific theory is true from a limited data sample, it is possible to prove that a theory is false. Suppose a scientist is considering the theory that all pieces of metal conduct electricity. Even if every piece of metal she examines does conduct electricity, this doesn’t prove that the theory is true, for reasons that we’ve seen. But if she finds even one piece of metal that does not conduct electricity, this does prove that the theory is false. For the inference from ‘this piece of metal does not conduct electricity’ to ‘it is false that all pieces of metal conduct electricity’ is a deductive inference – the premiss entails the conclusion. So if a scientist is only interested in demonstrating that a given theory is false, she may be able to accomplish her goal without the use of inductive inferences.

The weakness of Popper’s argument is obvious. For scientists are not only interested in showing that certain theories are false. When a scientist collects experimental data, her aim might be to show that a particular theory – her arch-rival’s theory perhaps – is false. But much more likely, she is trying to convince people that her own theory is true. And in order to do that, she will

have to resort to inductive reasoning of some sort. So Popper's attempt to show that science can get by without induction does not succeed.

Hume's problem

Although inductive reasoning is not logically watertight, it nonetheless seems like a perfectly sensible way of forming beliefs about the world. The fact that the sun has risen every day up until now may not prove that it will rise tomorrow, but surely it gives us very good reason to think it will? If you came across someone who professed to be entirely agnostic about whether the sun will rise tomorrow or not, you would regard them as very strange indeed, if not irrational.

But what justifies this faith we place in induction? How should we go about persuading someone who refuses to reason inductively that they are wrong? The 18th-century Scottish philosopher David Hume (1711–1776) gave a simple but radical answer to this question. He argued that the use of induction cannot be rationally justified at all. Hume admitted that we use induction all the time, in everyday life and in science, but he insisted this was just a matter of brute animal habit. If challenged to provide a good reason for using induction, we can give no satisfactory answer, he thought.

How did Hume arrive at this startling conclusion? He began by noting that whenever we make inductive inferences, we seem to presuppose what he called the 'uniformity of nature' (UN). To see what Hume means by this, recall some of the inductive inferences from the last section. We had the inference from 'my computer hasn't exploded up to now' to 'my computer won't explode today'; from 'all examined DS sufferers have an extra chromosome' to 'all DS sufferers have an extra chromosome'; from 'all bodies observed so far obey Newton's law of gravity' to 'all bodies obey Newton's law of gravity'; and so on. In each of these cases, our reasoning seems to depend on the assumption that objects we haven't examined will be similar, in the relevant respects, to objects of the same sort that we have examined. That assumption is what Hume means by the uniformity of nature.

But how do we know that the UN assumption is actually true, Hume asks? Can we perhaps prove its truth somehow (in the strict sense of proof)? No, says Hume, we cannot. For it is easy to imagine a universe where nature is not uniform, but changes its course randomly from day to day. In such a

universe, computers might sometimes explode for no reason, water might sometimes intoxicate us without warning, billiard balls might sometimes stop dead on colliding, and so on. Since such a 'non-uniform' universe is conceivable, it follows that we cannot strictly prove the truth of UN. For if we could prove that UN is true, then the non-uniform universe would be a logical impossibility.

Granted that we cannot prove UN, we might nonetheless hope to find good empirical evidence for its truth. After all, since UN has always held true up to now, surely that gives us good reason for thinking it is true? But this argument begs the question, says Hume! For it is itself an inductive argument, and so itself depends on the UN assumption. An argument that assumes UN from the outset clearly cannot be used to show that UN is true. To put the point another way, it is certainly an established fact that nature has behaved largely uniformly up to now. But we cannot appeal to this fact to argue that nature will continue to be uniform, because this assumes that what has happened in the past is a reliable guide to what will happen in the future – which *is* the uniformity of nature assumption. If we try to argue for UN on empirical grounds, we end up reasoning in a circle.

The force of Hume's point can be appreciated by imagining how you would go about persuading someone who doesn't trust inductive reasoning that they should. You would probably say: 'look, inductive reasoning has worked pretty well up until now. By using induction scientists have split the atom, landed men on the moon, invented computers, and so on. Whereas people who haven't used induction have tended to die nasty deaths. They have eaten arsenic believing that it would nourish them, jumped off tall buildings believing that they would fly, and so on (Figure 6). Therefore it will clearly pay you to reason inductively.' But of course this wouldn't convince the doubter. For to argue that induction is trustworthy because it has worked well up to now is to reason in an inductive way. Such an argument would carry no weight with someone who doesn't already trust induction. That is Hume's fundamental point.



6. What happens to people who don't trust induction.

So the position is this. Hume points out that our inductive inferences rest on the UN assumption. But we cannot prove that UN is true, and we cannot produce empirical evidence for its truth without begging the question. So our inductive inferences rest on an assumption about the world for which we have no good grounds. Hume concludes that our confidence in induction is just blind faith – it admits of no rational justification whatever.

This intriguing argument has exerted a powerful influence on the philosophy of science, and continues to do so today. (Popper's unsuccessful attempt to show that scientists need only use deductive inferences was motivated by his belief that Hume had shown the total irrationality of inductive reasoning.) The influence of Hume's argument is not hard to understand. For normally we think of science as the very paradigm of rational enquiry. We place great faith in what scientists tell us about the world. Every time we travel by aeroplane, we put our lives in the hands of the scientists who designed the plane. But science relies on induction, and Hume's argument seems to show that induction cannot be rationally justified. If Hume is right, the foundations on

which science is built do not look quite as solid as we might have hoped. This puzzling state of affairs is known as Hume's problem of induction.

Philosophers have responded to Hume's problem in literally dozens of different ways; this is still an active area of research today. Some people believe the key lies in the concept of probability. This suggestion is quite plausible. For it is natural to think that although the premisses of an inductive inference do not guarantee the truth of the conclusion, they do make it quite probable. So even if scientific knowledge cannot be certain, it may nonetheless be highly probable. But this response to Hume's problem generates difficulties of its own, and is by no means universally accepted; we will return to it in due course.

Another popular response is to admit that induction cannot be rationally justified, but to argue that this is not really so problematic after all. How might one defend such a position? Some philosophers have argued that induction is so fundamental to how we think and reason that it's not the sort of thing that could be justified. Peter Strawson, an influential contemporary philosopher, defended this view with the following analogy. If someone worried about whether a particular action was legal, they could consult the law-books and compare the action with what the law-books say. But suppose someone worried about whether the law itself was legal. This is an odd worry indeed. For the law is the standard against which the legality of other things is judged, and it makes little sense to enquire whether the standard itself is legal. The same applies to induction, Strawson argued. Induction is one of the standards we use to decide whether claims about the world are justified. For example, we use induction to judge whether a pharmaceutical company's claim about the amazing benefits of its new drug are justified. So it makes little sense to ask whether induction itself is justified.

Has Strawson really succeeded in defusing Hume's problem? Some philosophers say yes, others say no. But most people agree that it is very hard to see how there could be a satisfactory justification of induction. (Frank Ramsey, a Cambridge philosopher from the 1920s, said that to ask for a justification of induction was 'to cry for the moon'.) Whether this is something that should worry us, or shake our faith in science, is a difficult question that you should ponder for yourself.

Inference to the best explanation

The inductive inferences we've examined so far have all had essentially the same structure. In each case, the premiss of the inference has had the form 'all x's examined so far have been y', and the conclusion has had the form 'the next x to be examined will be y', or sometimes, 'all x's are y'. In other words, these inferences take us from examined to unexamined instances of a given kind.

Such inferences are widely used in everyday life and in science, as we have seen. However, there is another common type of non-deductive inference that doesn't fit this simple pattern. Consider the following example:

The cheese in the larder has disappeared, apart from a few crumbs
Scratching noises were heard coming from the larder last night

Therefore, the cheese was eaten by a mouse

It is obvious that this inference is non-deductive: the premisses do not entail the conclusion. For the cheese could have been stolen by the maid, who cleverly left a few crumbs to make it look like the handiwork of a mouse ([Figure 7](#)). And the scratching noises could have been caused in any number of ways – perhaps they were due to the boiler overheating. Nonetheless, the inference is clearly a reasonable one. For the hypothesis that a mouse ate the cheese seems to provide a better explanation of the data than do the various alternative explanations. After all, maids do not normally steal cheese, and modern boilers do not tend to overheat. Whereas mice do normally eat cheese when they get the chance, and do tend to make scratching sounds. So although we cannot be certain that the mouse hypothesis is true, on balance it looks quite plausible: it is the best way of accounting for the available data.



7. The mouse hypothesis and the maid hypothesis can both account for the missing cheese.

Reasoning of this sort is known as ‘inference to the best explanation’, for obvious reasons, or IBE for short. Certain terminological confusions surround the relation between IBE and induction. Some philosophers describe IBE as a type of inductive inference; in effect, they use ‘inductive inference’ to mean ‘any inference which is not deductive’. Others contrast IBE with inductive inference, as we have done above. On this way of cutting the pie, ‘inductive inference’ is reserved for inferences from examined to unexamined instances of a given kind, of the sort we examined earlier; IBE and inductive inference are then two different types of non-deductive inference. Nothing hangs on which choice of terminology we favour, so long as we stick to it consistently.

Scientists frequently use IBE. For example, Darwin argued for his theory of evolution by calling attention to various facts about the living world which are hard to explain if we assume that current species have been separately created, but which make perfect sense if current species have descended from common ancestors, as his theory held. For example, there are close anatomical similarities between the legs of horses and zebras. How do we explain this, if God created horses and zebras separately? Presumably he could have made their legs as different as he pleased. But if horses and

zebras have both descended from a recent common ancestor, this provides an obvious explanation of their anatomical similarity. Darwin argued that the ability of his theory to explain facts of this sort, and of many other sorts too, constituted strong evidence for its truth.

Another example of IBE is Einstein's famous work on Brownian motion. Brownian motion refers to the chaotic, zig-zag motion of microscopic particles suspended in a liquid or gas. It was discovered in 1827 by the Scottish botanist Robert Brown (1713–1858), while examining pollen grains floating in water. A number of attempted explanations of Brownian motion were advanced in the 19th century. One theory attributed the motion to electrical attraction between particles, another to agitation from external surroundings, and another to convection currents in the fluid. The correct explanation is based on the kinetic theory of matter, which says that liquids and gases are made up of atoms or molecules in motion. The suspended particles collide with the surrounding molecules, causing the erratic, random movements that Brown first observed. This theory was first proposed in the late 19th century but was not widely accepted, not least because many scientists didn't believe that atoms and molecules were real physical entities. But in 1905, Einstein provided an ingenious mathematical treatment of Brownian motion, making a number of precise, quantitative predictions which were later confirmed experimentally. After Einstein's work, the kinetic theory was quickly agreed to provide a far better explanation of Brownian motion than any of the alternatives, and scepticism about the existence of atoms and molecules rapidly subsided.

One interesting question is whether IBE or ordinary induction is a more fundamental pattern of inference. The philosopher Gilbert Harman has argued that IBE is more fundamental. According to this view, whenever we make an ordinary inductive inference such as 'all pieces of metal examined so far conduct electricity, therefore all pieces of metal conduct electricity' we are implicitly appealing to explanatory considerations. We assume that the correct explanation for why the pieces of metal in our sample conducted electricity, whatever it is, entails that all pieces of metal will conduct electricity; that is why we make the inductive inference. But if we believed, for example, that the explanation for why the pieces of metal in our sample conducted electricity was that a laboratory technician had tinkered with them, we would not infer that all pieces of metal conduct electricity. Proponents of this view do not say

there is no difference between IBE and ordinary induction – there clearly is. Rather, they think that ordinary induction is ultimately dependent on IBE.

However, other philosophers argue that this gets things backwards: IBE is itself parasitic on ordinary induction, they say. To see the grounds for this view, think back to the cheese-in-the-larder example above. Why do we regard the mouse hypothesis as a better explanation of the data than the maid hypothesis? Presumably, because we know that maids do not normally steal cheese, whereas mice do. But this is knowledge that we have gained through ordinary inductive reasoning, based on our previous observations of the behaviour of mice and maids. So according to this view, when we try to decide which of a group of competing hypotheses provides the best explanation of our data, we invariably appeal to knowledge that has been gained through ordinary induction. Thus it is incorrect to regard IBE as a more fundamental mode of inference.

Whichever of these opposing views we favour, one issue clearly demands more attention. If we want to use IBE, we need some way of deciding which of the competing hypotheses provides the best explanation of the data. But what criteria determine this? A popular answer is that the best explanation is the simplest or the most parsimonious one. Consider again the cheese-in-the-larder example. There are two pieces of data that need explaining: the missing cheese and the scratching noises. The mouse hypothesis postulates just one cause – a mouse – to explain both pieces of data. But the maid hypothesis must postulate two causes – a dishonest maid and an overheating boiler – to explain the same data. So the mouse hypothesis is more parsimonious, hence better. Similarly in the Darwin example. Darwin's theory could explain a very diverse range of facts about the living world, not just anatomical similarities between species. Each of these facts could be explained in other ways, as Darwin knew. But the theory of evolution explained all the facts in one go – that is what made it the best explanation of the data.

The idea that simplicity or parsimony is the mark of a good explanation is quite appealing, and certainly helps flesh out the idea of IBE. But if scientists use simplicity as a guide to inference, this raises a problem. For how do we know that the universe is simple rather than complex? Preferring a theory that explains the data in terms of the fewest number of causes does seem sensible. But is there any objective reason for thinking that such a theory is

more likely to be true than a less simple theory? Philosophers of science do not agree on the answer to this difficult question.

Probability and induction

The concept of probability is philosophically puzzling. Part of the puzzle is that the word 'probability' seems to have more than one meaning. If you read that the probability of an Englishwoman living to 100 years of age is 1 in 10, you would understand this as saying that one-tenth of all Englishwomen live to the age of 100. Similarly, if you read that the probability of a male smoker developing lung cancer is 1 in 4, you would take this to mean that a quarter of all male smokers develop lung cancer. This is known as the frequency interpretation of probability: it equates probabilities with proportions, or frequencies. But what if you read that the probability of finding life on Mars is 1 in 1,000? Does this mean that one out of every thousand planets in our solar system contains life? Clearly it does not. For one thing, there are only nine planets in our solar system. So a different notion of probability must be at work here.

One interpretation of the statement 'the probability of life on Mars is 1 in 1,000' is that the person who utters it is simply reporting a subjective fact about themselves – they are telling us how likely they think life on Mars is. This is the subjective interpretation of probability. It takes probability to be a measure of the strength of our personal opinions. Clearly, we hold some of our opinions more strongly than others. I am very confident that Brazil will win the World Cup, reasonably confident that Jesus Christ existed, and rather less confident that global environmental disaster can be averted. This could be expressed by saying that I assign a high probability to the statement 'Brazil will win the World Cup', a fairly high probability to 'Jesus Christ existed', and a low probability to 'global environmental disaster can be averted'. Of course, to put an exact number on the strength of my conviction in these statements would be hard, but advocates of the subjective interpretation regard this as a merely practical limitation. In principle, we should be able to assign a precise numerical probability to each of the statements about which we have an opinion, reflecting how strongly we believe or disbelieve them, they say.

The subjective interpretation of probability implies that there are no objective facts about probability, independently of what people believe. If I say that the probability of finding life on Mars is high and you say that it is very low, neither


of us is right or wrong – we are both simply stating how strongly we believe the statement in question. Of course, there is an objective fact about whether there is life on Mars or not; there is just no objective fact about how probable it is that there is life on Mars, according to the subjective interpretation.

The logical interpretation of probability rejects this position. It holds that a statement such as ‘the probability of life on Mars is high’ is objectively true or false, relative to a specified body of evidence. A statement’s probability is the measure of the strength of evidence in its favour, on this view. Advocates of the logical interpretation think that for any two statements in our language, we can in principle discover the probability of one, given the other as evidence. For example, we might want to discover the probability that there will be an ice age within 10,000 years, given the current rate of global warming. The subjective interpretation says there is no objective fact about this probability. But the logical interpretation insists that there is: the current rate of global warming confers a definite numerical probability on the occurrence of an ice age within 10,000 years, say 0.9 for example. A probability of 0.9 clearly counts as a high probability – for the maximum is 1 – so the statement ‘the probability that there will be an ice age within 10,000 years is high’ would then be objectively true, given the evidence about global warming.

If you have studied probability or statistics, you may be puzzled by this talk of different interpretations of probability. How do these interpretations tie in with what you learned? The answer is that the mathematical study of probability does not by itself tell us what probability means, which is what we have been examining above. Most statisticians would in fact favour the frequency interpretation, but the problem of how to interpret probability, like most philosophical problems, cannot be resolved mathematically. The mathematical formulae for working out probabilities remain the same, whichever interpretation we adopt.

This big yellow box is supposed to be here.

This big yellow box is supposed to be here.



Philosophers of science are interested in probability for two main reasons. The first is that in many branches of science, especially physics and biology, we find laws and theories that are formulated using the notion of probability. Consider, for example, the theory known as Mendelian genetics, which deals with the transmission of genes from one generation to another in sexually reproducing populations. One of the most important principles of Mendelian genetics is that every gene in an organism has a 50% chance of making it into any one of the organism's gametes (sperm or egg cells). Hence there is a 50% chance that any gene found in your mother will also be in you, and likewise for the genes in your father. Using this principle and others, geneticists can provide detailed explanations for why particular characteristics (e.g. eye colour) are distributed across the generations of a family in the way that they are. Now 'chance' is just another word for probability, so it is obvious that our Mendelian principle makes essential use of the concept of probability.

Many other examples could be given of scientific laws and principles that are expressed in terms of probability. The need to understand these laws and principles is an important motivation for the philosophical study of probability.

The second reason why philosophers of science are interested in the concept of probability is the hope that it might shed some light on inductive inference, in particular on Hume's problem; this shall be our focus here. At the root of Hume's problem is the fact that the premisses of an inductive inference do not guarantee the truth of its conclusion. But it is tempting to suggest that the premisses of a typical inductive inference do make the conclusion highly probable. Although the fact that all objects examined so far obey Newton's law of gravity doesn't prove that all objects do, surely it does make it very probable? So surely Hume's problem can be answered quite easily after all?

However, matters are not quite so simple. For we must ask what interpretation of probability this response to Hume assumes. On the frequency interpretation, to say it is highly probable that all objects obey Newton's law is to say that a very high proportion of all objects obey the law. But there is no way we can know that, unless we use induction! For we have only examined a tiny fraction of all the objects in the universe. So Hume's problem remains. Another way to see the point is this. We began with the inference from 'all examined objects obey Newton's law' to 'all objects obey Newton's law'. In response to Hume's worry that the premiss of this inference doesn't guarantee the truth of the conclusion, we suggested that it might nonetheless make the conclusion highly probable. But the inference from 'all examined objects obey Newton's law' to 'it is highly probable that all objects obey Newton's law' is still an inductive inference, given that the latter means 'a very high proportion of all objects obey Newton's law', as it does according to the frequency interpretation. So appealing to the concept of probability does not take the sting out of Hume's argument, if we adopt a frequency interpretation of probability. For knowledge of probabilities then becomes itself dependent on induction.

The subjective interpretation of probability is also powerless to solve Hume's problem, though for a different reason. Suppose John believes that the sun will rise tomorrow and Jack believes it will not. They both accept the evidence that the sun has risen every day in the past. Intuitively, we want to say that John is rational and Jack isn't, because the evidence makes John's belief more probable. But if probability is simply a matter of subjective opinion, we

cannot say this. All we can say is that John assigns a high probability to 'the sun will rise tomorrow' and Jack does not. If there are no objective facts about probability, then we cannot say that the conclusions of inductive inferences are objectively probable. So we have no explanation of why someone like Jack, who declines to use induction, is irrational. But Hume's problem is precisely the demand for such an explanation.

The logical interpretation of probability holds more promise of a satisfactory response to Hume. Suppose there is an objective fact about the probability that the sun will rise tomorrow, given that it has risen every day in the past. Suppose this probability is very high. Then we have an explanation of why John is rational and Jack isn't. For John and Jack both accept the evidence that the sun has risen every day in the past, but Jack fails to realize that this evidence makes it highly probable that the sun will rise tomorrow, while John does realize this. Regarding a statement's probability as a measure of the evidence in its favour, as the logical interpretation recommends, tallies neatly with our intuitive feeling that the premisses of an inductive inference can make the conclusion highly probable, even if they cannot guarantee its truth.

Unsurprisingly, therefore, those philosophers who have tried to solve Hume's problem via the concept of probability have tended to favour the logical interpretation. (One of these was the famous economist John Maynard Keynes, whose early interests were in logic and philosophy.) Unfortunately, most people today believe that the logical interpretation of probability faces very serious, probably insuperable, difficulties. This is because all the attempts to work out the logical interpretation of probability in any detail have run up against a host of problems, both mathematical and philosophical. As a result, many philosophers today are inclined to reject outright the underlying assumption of the logical interpretation – that there are objective facts about the probability of one statement, given another. Rejecting this assumption leads naturally to the subjective interpretation of probability, but that, as we have seen, offers scant hope of a satisfactory response to Hume.

Even if Hume's problem is ultimately insoluble, as seems likely, thinking about the problem is still a valuable exercise. For reflecting on the problem of induction leads us into a thicket of interesting questions about the structure of scientific reasoning, the nature of rationality, the appropriate degree of confidence to place in science, the interpretation of probability, and more. Like most philosophical questions, these questions probably do not admit of final

answers, but in grappling with them we learn much about the nature and limits of scientific knowledge.