

22/07/05
X 3165P
Rs. 175/-

Samir Okasha

PHILOSOPHY OF SCIENCE

501
Oka

A Very Short Introduction

14046

501
Oka



14046

OXFORD
UNIVERSITY PRESS

OXFORD

UNIVERSITY PRESS

Great Clarendon Street, Oxford o x 2 6 d p

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide in

Oxford New York

Auckland Bangkok Buenos Aires Cape Town Chennai
Dar es Salaam Delhi Hong Kong Istanbul Karachi Kolkata
Kuala Lumpur Madrid Melbourne Mexico City Mumbai Nairobi
São Paulo Shanghai Taipei Tokyo Toronto

Oxford is a registered trade mark of Oxford University Press
in the UK and in certain other countries

Published in the United States
by Oxford University Press Inc., New York

© Samir Okasha 2002

The moral rights of the author have been asserted

Database right Oxford University Press (maker)

First published as a Very Short Introduction 2002

All rights reserved. No part of this publication may be reproduced,
stored in a retrieval system, or transmitted, in any form or by any means,
without the prior permission in writing of Oxford University Press,
or as expressly permitted by law, or under terms agreed with the appropriate
reprographics rights organizations. Enquiries concerning reproduction
outside the scope of the above should be sent to the Rights Department,
Oxford University Press, at the address above

You must not circulate this book in any other binding or cover
and you must impose this same condition on any acquirer

British Library Cataloguing in Publication Data

Data available

Library of Congress Cataloging in Publication Data

Data available

ISBN 0-19-280283-6

5 7 9 10 8 6 4

Typeset by RefineCatch Ltd, Bungay, Suffolk

Printed in Great Britain by

TJ International Ltd., Padstow, Cornwall

Chapter 3

Explanation in science

One of the most important aims of science is to try and explain what happens in the world around us. Sometimes we seek explanations for practical ends. For example, we might want to know why the ozone layer is being depleted so quickly, in order to try and do something about it. In other cases we seek scientific explanations simply to satisfy our intellectual curiosity – we want to understand more about how the world works. Historically, the pursuit of scientific explanation has been motivated by both goals.

Quite often, modern science is successful in its aim of supplying explanations. For example, chemists can explain why sodium turns yellow when it burns. Astronomers can explain why solar eclipses occur when they do. Economists can explain why the yen declined in value in the 1980s. Geneticists can explain why male baldness tends to run in families. Neurophysiologists can explain why extreme oxygen deprivation leads to brain damage. You can probably think of many other examples of successful scientific explanations.

But what exactly *is* scientific explanation? What exactly does it mean to say that a phenomenon can be ‘explained’ by science? This is a question that has exercised philosophers since Aristotle, but our starting point will be a famous account of scientific explanation put forward in the 1950s by the American philosopher Carl Hempel.

Hempel's account is known as the *covering law* model of explanation, for reasons that will become clear.

Hempel's covering law model of explanation

The basic idea behind the covering law model is straightforward. Hempel noted that scientific explanations are usually given in response to what he called 'explanation-seeking why questions'. These are questions such as 'why is the earth not perfectly spherical?', 'why do women live longer than men?', and the like – they are demands for explanation. To give a scientific explanation is thus to provide a satisfactory answer to an explanation-seeking why question. If we could determine the essential features that such an answer must have, we would know what scientific explanation is.

Hempel suggested that scientific explanations typically have the logical structure of an argument, i.e. a set of premisses followed by a conclusion. The conclusion states that the phenomenon that needs explaining actually occurs, and the premisses tell us why the conclusion is true. Thus suppose someone asks why sugar dissolves in water. This is an explanation-seeking why question. To answer it, says Hempel, we must construct an argument whose conclusion is 'sugar dissolves in water' and whose premisses tell us why this conclusion is true. The task of providing an account of scientific explanation then becomes the task of characterizing precisely the relation that must hold between a set of premisses and a conclusion, in order for the former to count as an explanation of the latter. That was the problem Hempel set himself.

Explanation in science

Hempel's answer to the problem was three-fold. Firstly, the premisses should entail the conclusion, i.e. the argument should be a deductive one. Secondly, the premisses should all be true. Thirdly, the premisses should consist of at least one general law. General laws are things such as 'all metals conduct electricity', 'a body's acceleration varies inversely with its mass', 'all plants contain chlorophyll', and so on; they contrast with particular facts such as

'this piece of metal conducts electricity', 'the plant on my desk contains chlorophyll' and so on. General laws are sometimes called 'laws of nature'. Hempel allowed that a scientific explanation could appeal to particular facts as well as general laws, but he held that at least one general law was always essential. So to explain a phenomenon, on Hempel's conception, is to show that its occurrence follows deductively from a general law, perhaps supplemented by other laws and/or particular facts, all of which must be true.

To illustrate, suppose I am trying to explain why the plant on my desk has died. I might offer the following explanation. Owing to the poor light in my study, no sunlight has been reaching the plant; but sunlight is necessary for a plant to photosynthesize; and without photosynthesis a plant cannot make the carbohydrates it needs to survive, and so will die; therefore my plant died. This explanation fits Hempel's model exactly. It explains the death of the plant by deducing it from two true laws – that sunlight is necessary for photosynthesis, and that photosynthesis is necessary for survival – and one particular fact – that the plant was not getting any sunlight. Given the truth of the two laws and the particular fact, the death of the plant *had* to occur; that is why the former constitute a good explanation of the latter.

Schematically, Hempel's model of explanation can be written as follows:

General laws

Particular facts

⇒

Phenomenon to be explained

The phenomenon to be explained is called the *explanandum*, and the general laws and particular facts that do the explaining are called the *explanans*. The explanandum itself may be either a particular fact or a general law. In the example above, it was a particular fact – the death of my plant. But sometimes the things we

want to explain are general. For example, we might wish to explain why exposure to the sun leads to skin cancer. This is a general law, not a particular fact. To explain it, we would need to deduce it from still more fundamental laws – presumably, laws about the impact of radiation on skin cells, combined with particular facts about the amount of radiation in sunlight. So the structure of a scientific explanation is essentially the same, whether the *explanandum*, i.e. the thing we are trying to explain, is particular or general.

It is easy to see why Hempel's model is called the covering law model of explanation. For according to the model, the essence of explanation is to show that the phenomenon to be explained is 'covered' by some general law of nature. There is certainly something appealing about this idea. For showing that a phenomenon is a consequence of a general law does in a sense take the mystery out of it – it renders it more intelligible. And in fact, scientific explanations do often fit the pattern Hempel describes. For example, Newton explained why the planets move in ellipses around the sun by showing that this can be deduced from his law of universal gravitation, along with some minor additional assumptions. Newton's explanation fits Hempel's model exactly: a phenomenon is explained by showing that it had to be so, given the laws of nature plus some additional facts. After Newton, there was no longer any mystery about why planetary orbits are elliptical.

Hempel was aware that not all scientific explanations fit his model exactly. For example, if you ask someone why Athens is always immersed in smog, they will probably say 'because of car exhaust pollution'. This is a perfectly acceptable scientific explanation, though it involves no mention of any laws. But Hempel would say that if the explanation were spelled out in full detail, laws would enter the picture. Presumably there is a law that says something like 'if carbon monoxide is released into the earth's atmosphere in sufficient concentration, smog clouds will form'. The full explanation of why Athens is bathed in smog would cite this law, along with the fact that car exhaust contains carbon monoxide and

Athens has lots of cars. In practice, we wouldn't spell out the explanation in this much detail unless we were being very pedantic. But if we were to spell it out, it would correspond quite well to the covering law pattern.


Hempel drew an interesting philosophical consequence from his model about the relation between explanation and prediction. He argued that these are two sides of the same coin. Whenever we give a covering law explanation of a phenomenon, the laws and particular facts we cite would have enabled us to predict the occurrence of the phenomenon, if we hadn't already known about it. To illustrate, consider again Newton's explanation of why planetary orbits are elliptical. This fact was known long before Newton explained it using his theory of gravity – it was discovered by Kepler. But if it had not been known, Newton would have been able to predict it from his theory of gravity, for his theory entails that planetary orbits are elliptical, given minor additional assumptions. Hempel expressed this by saying that every scientific explanation is potentially a prediction – it would have served to predict the phenomenon in question, had it not already been known. The converse was also true, Hempel thought: every reliable prediction is potentially an explanation. To illustrate, suppose scientists predict that mountain gorillas will be extinct by 2010, based on information about the destruction of their habitat. Suppose they turn out to be right. According to Hempel, the information they used to predict the gorillas' extinction before it happened will serve to explain that same fact after it has happened. Explanation and prediction are structurally symmetric.

Though the covering law model captures the structure of many actual scientific explanations quite well, it also faces a number of awkward counter-examples. These counter-examples fall into two classes. On the one hand, there are cases of genuine scientific explanations that do not fit the covering law model, even approximately. These cases suggest that Hempel's model is too strict – it excludes some *bona fide* scientific explanations. On the

other hand, there are cases of things that *do* fit the covering law model, but intuitively do not count as genuine scientific explanations. These cases suggest that Hempel's model is too liberal – it allows in things that should be excluded. We will focus on counter-examples of the second sort.

The problem of symmetry

Suppose you are lying on the beach on a sunny day, and you notice that a flagpole is casting a shadow of 20 metres across the sand (Figure 8).



8. A 15-metre flagpole casts a shadow of 20 metres on the beach when the sun is 37° overhead.

Someone asks you to explain why the shadow is 20 metres long. This is an explanation-seeking why question. A plausible answer might go as follows: 'light rays from the sun are hitting the flagpole, which is exactly 15 metres high. The angle of elevation of the sun is 37° . Since light travels in straight lines, a simple trigonometric calculation ($\tan 37^\circ = 15/20$) shows that the flagpole will cast a shadow 20 metres long'.

This looks like a perfectly good scientific explanation. And by rewriting it in accordance with Hempel's schema, we can see that it fits the covering law model:

General laws	Light travels in straight lines Laws of trigonometry
Particular facts	Angle of elevation of the sun is 37° Flagpole is 15 metres high
$\Rightarrow$	
Phenomenon to be explained	Shadow is 20 metres long

The length of the shadow is deduced from the height of the flagpole and the angle of elevation of the sun, along with the optical law that light travels in straight lines and the laws of trigonometry. Since these laws are true, and since the flagpole is indeed 15 metres high, the explanation satisfies Hempel's requirements precisely. So far so good. The problem arises as follows. Suppose we swap the *explanandum* – that the shadow is 20 metres long – with the particular fact that the flagpole is 15 metres high. The result is this:

General law	Light travels in straight lines Laws of trigonometry
Particular facts	Angle of elevation of the sun is 37° Shadow is 20 metres long
$\Rightarrow$	
Phenomenon to be explained	Flagpole is 15 metres high

This 'explanation' clearly conforms to the covering law pattern too. The height of the flagpole is deduced from the length of the shadow it casts and the angle of elevation of the sun, along with the optical law that light travels in straight lines and the laws of trigonometry. But it seems very odd to regard this as an *explanation* of why the flagpole is 15 metres high. The real explanation of why the flagpole is 15 metres high is presumably that a carpenter deliberately made it so – it has nothing to do with the length of the shadow that it casts. So Hempel's model is too liberal: it allows something to count as a scientific explanation that obviously is not.

The general moral of the flagpole example is that the concept of explanation exhibits an important asymmetry. The height of the flagpole explains the length of the shadow, given the relevant laws and additional facts, but not vice-versa. In general, if x explains y , given the relevant laws and additional facts, then it will not be true that y explains x , given the same laws and facts. This is sometimes expressed by saying that explanation is an asymmetric relation. Hempel's covering law model does not respect this asymmetry. For just as we can deduce the length of the shadow from the height of the flagpole, given the laws and additional facts, so we can deduce the height of the flagpole from the length of the shadow. In other words, the covering law model implies that explanation should be a symmetric relation, but in fact it is asymmetric. So Hempel's model fails to capture fully what it is to be a scientific explanation.

The shadow and flagpole case also provides a counter-example to Hempel's thesis that explanation and prediction are two sides of the same coin. The reason is obvious. Suppose you didn't know how high the flagpole was. If someone told you that it was casting a shadow of 20 metres and that the sun was 37° overhead, you would be able to *predict* the flagpole's height, given that you knew the relevant optical and trigonometrical laws. But as we have just seen, this information clearly doesn't *explain* why the flagpole has the height it does. So in this example prediction and explanation part ways. Information that serves to predict a fact before we know it does not serve to explain that same fact after we know it, which contradicts Hempel's thesis.

The problem of irrelevance

Suppose a young child is in a hospital in a room full of pregnant women. The child notices that one person in the room – who is a man called John – is not pregnant, and asks the doctor why not. The doctor replies: 'John has been taking birth-control pills regularly for the last few years. People who take birth-control pills regularly never become pregnant. Therefore, John has not become pregnant'.

Let us suppose for the sake of argument that what the doctor says is true – John is mentally ill and does indeed take birth-control pills, which he believes help him. Even so, the doctor's reply to the child is clearly not very helpful. The correct explanation of why John has not become pregnant, obviously, is that he is male and males cannot become pregnant.

However, the explanation the doctor has given the child fits the covering law model perfectly. The doctor deduces the phenomenon to be explained – that John is not pregnant – from the general law that people who take birth-control pills do not become pregnant and the particular fact that John has been taking birth-control pills. Since both the general law and the particular fact are true, and since they do indeed entail the *explanandum*, according to the covering law model the doctor has given a perfectly adequate explanation of why John is not pregnant. But of course he hasn't. Hence the covering law model is again too permissive: it allows things to count as scientific explanations that intuitively are not.

The general moral is that a good explanation of a phenomenon should contain information that is *relevant* to the phenomenon's occurrence. This is where the doctor's reply to the child goes wrong. Although what the doctor tells the child is perfectly true, the fact that John has been taking birth-control pills is irrelevant to his not being pregnant, because he wouldn't have been pregnant even if he hadn't been taking the pills. This is why the doctor's reply does not constitute a good answer to the child's question. Hempel's model does not respect this crucial feature of our concept of explanation.

Explanation and causality

Since the covering law model encounters so many problems, it is natural to look for an alternative way of understanding scientific explanation. Some philosophers believe that the key lies in the concept of causality. This is quite an attractive suggestion. For in many cases to explain a phenomenon is indeed to say what caused

it. For example, if an accident investigator is trying to explain an aeroplane crash, he is obviously looking for the cause of the crash. Indeed, the questions 'why did the plane crash?' and 'what was the cause of the plane crash?' are practically synonymous. Similarly, if an ecologist is trying to explain why there is less biodiversity in the tropical rainforests than there used to be, he is clearly looking for the cause of the reduction in biodiversity. The link between the concepts of explanation and causality is quite intimate.

Impressed by this link, a number of philosophers have abandoned the covering law account of explanation in favour of causality-based accounts. The details vary, but the basic idea behind these accounts is that to explain a phenomenon is simply to say what caused it. In some cases, the difference between the covering law and causal accounts is not actually very great, for to deduce the occurrence of a phenomenon from a general law often just *is* to give its cause. For example, recall again Newton's explanation of why planetary orbits are elliptical. We saw that this explanation fits the covering law model – for Newton deduced the shape of the planetary orbits from his law of gravity, plus some additional facts. But Newton's explanation was also a causal one, since elliptical planetary orbits are caused by the gravitational attraction between planets and the sun.

However, the covering law and causal accounts are not fully equivalent – in some cases they diverge. Indeed, many philosophers favour a causal account of explanation precisely because they think it can avoid some of the problems facing the covering law model. Recall the flagpole problem. Why do our intuitions tell us that the height of the flagpole explains the length of the shadow, given the laws, but not vice-versa? Plausibly, because the height of the flagpole is the cause of the shadow being 20 metres long, but the shadow being 20 metres long is not the cause of the flagpole being 15 metres high. So unlike the covering law model, a causal account of explanation gives the 'right' answer in the flagpole case – it respects our intuition that we cannot

explain the height of the flagpole by pointing to the length of the shadow it casts.

The general moral of the flagpole problem was that the covering law model cannot accommodate the fact that explanation is an asymmetric relation. Now causality is obviously an asymmetric relation too: if x is the cause of y , then y is not the cause of x . For example, if the short-circuit caused the fire, then the fire clearly did not cause the short-circuit. It is therefore quite plausible to suggest that the asymmetry of explanation derives from the asymmetry of causality. If to explain a phenomenon is to say what caused it, then since causality is asymmetric we should expect explanation to be asymmetric too – as it is. The covering law model runs up against the flagpole problem precisely because it tries to analyse the concept of scientific explanation without reference to causality.

The same is true of the birth-control pill case. That John takes birth-control pills does not explain why he isn't pregnant, because the birth-control pills are not the cause of his not being pregnant. Rather, John's gender is the cause of his not being pregnant. That is why we think that the correct answer to the question 'why is John not pregnant?' is 'because he is a man, and men can't become pregnant', rather than the doctor's answer. The doctor's answer satisfies the covering law model, but since it does not correctly identify the cause of the phenomenon we wish to explain, it does not constitute a genuine explanation. The general moral we drew from the birth-control pill example was that a genuine scientific explanation must contain information that is relevant to the *explanandum*. In effect, this is another way of saying that the explanation should tell us the *explanandum*'s cause. Causality-based accounts of scientific explanation do not run up against the problem of irrelevance.

It is easy to criticize Hempel for failing to respect the close link between causality and explanation, and many people have done so.

In some ways, this criticism is a bit unfair. For Hempel subscribed to a philosophical doctrine known as *empiricism*, and empiricists are traditionally very suspicious of the concept of causality. Empiricism says that all our knowledge comes from experience. David Hume, whom we met in the last chapter, was a leading empiricist, and he argued that it is impossible to experience causal relations. So he concluded that they don't exist – causality is a figment of our imagination! This is a very hard conclusion to accept. Surely it is an objective fact that dropping glass vases causes them to break? Hume denied this. He allowed that it is an objective fact that most glass vases that have been dropped have in fact broken. But our idea of causality includes more than this. It includes the idea of a causal link between the dropping and the breaking, i.e. that the former brings about the latter. No such links are to be found in the world, according to Hume: all we see is a vase being dropped, and then it breaking a moment later. We experience no causal connection between the first event and the second. Causality is therefore a fiction.

Most empiricists have not accepted this startling conclusion outright. But as a result of Hume's work, they have tended to regard causality as a concept to be treated with great caution. So to an empiricist, the idea of analysing the concept of explanation in terms of the concept of causality would seem perverse. If one's goal is to clarify the concept of scientific explanation, as Hempel's was, there is little point in using notions that are equally in need of clarification themselves. And for empiricists, causality is definitely in need of philosophical clarification. So the fact that the covering law model makes no mention of causality was not a mere oversight on Hempel's part. In recent years, empiricism has declined somewhat in popularity. Furthermore, many philosophers have come to the conclusion that the concept of causality, although philosophically problematic, is indispensable to how we understand the world. So the idea of a causality-based account of scientific explanation seems more acceptable than it would have done in Hempel's day.

Causality-based accounts of explanation certainly capture the structure of many actual scientific explanations quite well, but are they the whole story? Many philosophers say no, on the grounds that certain scientific explanations do not seem to be causal. One type of example stems from what are called 'theoretical identifications' in science. Theoretical identifications involve identifying one concept with another, usually drawn from a different branch of science. 'Water is H_2O ' is an example, as is 'temperature is average molecular kinetic energy'. In both of these cases, a familiar everyday concept is equated or identified with a more esoteric scientific concept. Often, theoretical identifications furnish us with what seem to be scientific explanations. When chemists discovered that water is H_2O , they thereby explained what water is. Similarly, when physicists discovered that an object's temperature is the average kinetic energy of its molecules, they thereby explained what temperature is. But neither of these explanations is causal. Being made of H_2O doesn't *cause* a substance to be water – it just *is* being water. Having a particular average molecular kinetic energy doesn't *cause* a liquid to have the temperature it does – it just *is* having that temperature. If these examples are accepted as legitimate scientific explanations, they suggest that causality-based accounts of explanation cannot be the whole story.

Can science explain everything?

Modern science can explain a great deal about the world we live in. But there are also numerous facts that have not been explained by science, or at least not explained fully. The origin of life is one such example. We know that about 4 billion years ago, molecules with the ability to make copies of themselves appeared in the primeval soup, and life evolved from there. But we do not understand how these self-replicating molecules got there in the first place. Another example is the fact that autistic children tend to have very good memories. Numerous studies of autistic children have confirmed this fact, but as yet nobody has succeeded in explaining it.

Many people believe that in the end, science will be able to explain facts of this sort. This is quite a plausible view. Molecular biologists are working hard on the problem of the origin of life, and only a pessimist would say they will never solve it. Admittedly, the problem is not easy, not least because it is very hard to know what conditions on earth 4 billion years ago were like. But nonetheless, there is no reason to think that the origin of life will never be explained. Similarly for the exceptional memories of autistic children. The science of memory is still in its infancy, and much remains to be discovered about the neurological basis of autism. Obviously we cannot guarantee that the explanation will eventually be found. But given the number of explanatory successes that modern science has already notched up, the smart money must be on many of today's unexplained facts eventually being explained too.

But does this mean that science can in principle explain everything? Or are there some phenomena that must forever elude scientific explanation? This is not an easy question to answer. On the one hand, it seems arrogant to assert that science can explain everything. On the other hand, it seems short-sighted to assert that any particular phenomenon can never be explained scientifically. For science changes and develops very fast, and a phenomenon that looks completely inexplicable from the vantage-point of today's science may be easily explained tomorrow.

According to some philosophers, there is a purely logical reason why science will never be able to explain everything. For in order to explain something, whatever it is, we need to invoke something else. But what explains the second thing? To illustrate, recall that Newton explained a diverse range of phenomena using his law of gravity. But what explains the law of gravity itself? If someone asks *why* all bodies exert a gravitational force on each other, what should we tell them? Newton had no answer to this question. In Newtonian science the law of gravity was a fundamental principle: it explained other things, but could not itself be explained. The

moral is generalizable. However much the science of the future can explain, the explanations it gives will have to make use of certain fundamental laws and principles. Since nothing can explain itself, it follows that at least some of these laws and principles will themselves remain unexplained.

Whatever one makes of this argument, it is undeniably very abstract. It purports to show that some things will never be explained, but does not tell us what they are. However, some philosophers have made concrete suggestions about phenomena that they think science can never explain. An example is consciousness – the distinguishing feature of thinking, feeling creatures such as ourselves and other higher animals. Much research into the nature of consciousness has been and continues to be done, by brain scientists, psychologists, and others. But a number of recent philosophers claim that whatever this research throws up, it will never fully explain the nature of consciousness. There is something intrinsically mysterious about the phenomenon of consciousness, they maintain, that no amount of scientific investigation can eliminate.

What are the grounds for this view? The basic argument is that conscious experiences are fundamentally unlike anything else in the world, in that they have a 'subjective aspect'. Consider, for example, the experience of watching a terrifying horror movie. This is an experience with a very distinctive 'feel' to it; in the current jargon, there is 'something that it is like' to have the experience. Neuroscientists may one day be able to give a detailed account of the complex goings-on in the brain that produce our feeling of terror. But will this explain why watching a horror movie feels the way it does, rather than feeling some other way? Many people believe that it will not. On this view, the scientific study of the brain can at most tell us which brain processes are correlated with which conscious experiences. This is certainly interesting and valuable information. However, it doesn't tell us *why* experiences with distinctive subjective 'feels' should result from the purely physical

goings-on in the brain. Hence consciousness, or at least one important aspect of it, is scientifically inexplicable.

Though quite compelling, this argument is very controversial and not endorsed by all philosophers, let alone all neuroscientists. Indeed, a well-known book published in 1991 by the philosopher Daniel Dennett is defiantly entitled *Consciousness Explained*. Supporters of the view that consciousness is scientifically inexplicable are sometimes accused of having a lack of imagination. Even if it is true that brain science as currently practised cannot explain the subjective aspect of conscious experience, can we not imagine the emergence of a radically different type of brain science, with radically different explanatory techniques, that *does* explain why our experiences feel the way they do? There is a long tradition of philosophers trying to tell scientists what is and isn't possible, and later scientific developments have often proved the philosophers wrong. Only time will tell whether the same fate awaits those who argue that consciousness must always elude scientific explanation.

Explanation and reduction

The different scientific disciplines are designed for explaining different types of phenomena. To explain why rubber doesn't conduct electricity is a task for physics. To explain why turtles have such long lives is a task for biology. To explain why higher interest rates reduce inflation is a task for economics, and so on. In short, there is a division of labour between the different sciences: each specializes in explaining its own particular set of phenomena. This explains why the sciences are not usually in competition with one another – why biologists, for example, do not worry that physicists and economists might encroach on their turf.

Nonetheless, it is widely held that the different branches of science are not all on a par: some are more fundamental than others. Physics is usually regarded as the most fundamental science of all.

Why? Because the objects studied by the other sciences are ultimately composed of physical particles. Consider living organisms, for example. Living organisms are made up of cells, which are themselves made up of water, nucleic acids (such as DNA), proteins, sugars, and lipids (fats), all of which consist of molecules or long chains of molecules joined together. But molecules are made up of atoms, which are physical particles. So the objects biologists study are ultimately just very complex physical entities. The same applies to the other sciences, even the social sciences. Take economics, for example. Economics studies the behaviour of corporations and consumers in the market place, and the consequences of this behaviour. But consumers are human beings and corporations are made up of human beings; and human beings are living organisms, hence physical entities.

Does this mean that, in principle, physics can subsume all the higher-level sciences? Since everything is made up of physical particles, surely if we had a complete physics, which allowed us to predict perfectly the behaviour of every physical particle in the universe, all the other sciences would become superfluous? Most philosophers resist this line of thought. After all, it seems crazy to suggest that physics might one day be able to explain the things that biology and economics explain. The prospect of deducing the laws of biology and economics straight from the laws of physics looks very remote. Whatever the physics of the future looks like, it is most unlikely to be capable of predicting economic downturns. Far from being reducible to physics, sciences such as biology and economics seem largely autonomous of it.

This leads to a philosophical puzzle. How can a science that studies entities that are ultimately physical *not* be reducible to physics? Granted that the higher-level sciences are in fact autonomous of physics, how is this possible? According to some philosophers, the answer lies in the fact that the objects studied by the higher-level sciences are 'multiply realized' at the physical level. To illustrate the idea of multiple realization, imagine a collection of ashtrays. Each

individual ashtray is obviously a physical entity, like everything else in the universe. But the physical composition of the ashtrays could be very different – some might be made of glass, others of aluminium, others of plastic, and so on. And they will probably differ in size, shape, and weight. There is virtually no limit on the range of different physical properties that an ashtray can have. So it is impossible to define the concept ‘ashtray’ in purely physical terms. We cannot find a true statement of the form ‘ x is an ashtray if and only if x is’ where the blank is filled by an expression taken from the language of physics. This means that ashtrays are multiply realized at the physical level.

Philosophers have often invoked multiple realization to explain why psychology cannot be reduced to physics or chemistry, but in principle the explanation works for any higher-level science. Consider, for example, the biological fact that nerve cells live longer than skin cells. Cells are physical entities, so one might think that this fact will one day be explained by physics. However, cells are almost certainly multiply realized at the microphysical level. Cells are ultimately made up of atoms, but the precise arrangement of atoms will be very different in different cells. So the concept ‘cell’ cannot be defined in terms drawn from fundamental physics. There is no true statement of the form ‘ x is a cell if and only if x is’ where the blank is filled by an expression taken from the language of microphysics. If this is correct, it means that fundamental physics will never be able to explain why nerve cells live longer than skin cells, or indeed any other facts about cells. The vocabulary of cell biology and the vocabulary of physics do not map onto each other in the required way. Thus we have an explanation of why it is that cell biology cannot be reduced to physics, despite the fact that cells are physical entities. Not all philosophers are happy with the doctrine of multiple realization, but it does promise to provide a neat explanation of the autonomy of the higher-level sciences, both from physics and from each other.