What's wrong with Reductionism?

What is Reductionism?

For a scientist, one of the pleasures of life is taking things to pieces. There are some tedious occasions when a little spring jumps out of the works and is lost for ever in the carpet, but taking something to pieces and putting it together again is part of a scientist's life because it gives an insight into how things work. You can drive a car or ride a bicycle much better when you have taken them to pieces because you know from experience what each knob or lever is going to do. Regrettably, in these days machines are controlled by sealed black boxes, whose functions you can only guess at. Modern cars are controlled by computers, which makes it difficult to know what is going on, compared with the old days when a Rolls-Royce had a big lever on the steering wheel to alter the engine timing. A driver in those days was much more aware of what was going on under the bonnet than we are now.

Perhaps it is because modern life is so complicated, that taking things to pieces is rather frowned on these days. Reductionism, that is understanding the whole by understanding the elements that make it up, is a fundamental scientific method, but in the process, scientists are often accused of oversimplifying complex situations. This is a criticism that is sometimes justified: people are not just walking bags of chemicals in spite of what some biochemists might say. But does something mysterious happen when you assemble a whole from its parts? Do you destroy something when you take it to pieces? Is science inevitably limited because it proceeds by reductionism?

I would say not. In fact I will go further and say that it is not possible to think about the real world without employing reductionism in some form or another. But you need to know what it is you are doing and what are the limitations of the method of reductionism you are using. As an example, if you were interested in the optical properties of water, you would need to know how light interacts with individual water molecules. These days, this is something you can calculate and so by focussing on one individual molecule you can explain why water is colourless. But focussing on one molecule means you have lost touch with whether this is a molecule of water, ice or steam and thereby have lost sight of a whole host of effects which make up the fascinating properties of water, and which, incidentally, make life possible. It is interesting that the physical state of a substance, whether it is a solid, liquid or a gas, simply does not apply when you are considering individual molecules. For a molecule on its own, *solid*, *liquid* or *gas* are meaningless terms. It is not that something magic happens when you bring individual molecules of water together to make a raindrop or a snowflake, but looking at one molecule on its own will not tell you how it happens.

The science of liquids, solids and gases is of course very well known, so the scientific method, which includes reductionism, can apply to collective behaviour. In biology we have the science of ecology, concerned with the interaction between plants and animals and the environment and cosmology is concerned with the evolution and interaction of stars and galaxies, both very large scale collections of objects In all science there is some domain which the science is applied to, which may be minute or immense. For cosmology this domain is the whole universe, but for a particle physicist it could be an individual proton or electron. But in order to reason about what is going on within the domain, scientists make simplifications. The simple kinetic theory of gases for example, treats the molecules making up the gas as hard, round perfectly elastic spheres, rather than complicated structures of electrons and protons. There is no real limit to the scale of a domain of investigation, but the larger the scale, the more one simplifies. How else are you to get your head around a problem?

So scientific thinking is characterised by a domain of application and a simplification within that domain and this is what constitutes reductionism. This structure, required because of the complexity and extent of reality compared with our ability to think, means that it applies to any thinking about reality. If you are not doing reductionism, you are not thinking about reality. You might be doing mathematics or music or some other activity, but you are not thinking about reality. The question is, what about religion? How much of that is to do with reality and how much to do with spiritual exercise as a guide for life?

The bit of reality which religion is concerned with is ourselves: our hopes and fears, our values and goals. It is hard to think of a more complicated bit of reality to deal with, so when using reductionism, as use it I feel we must, it is important to be clear about what this process is and what are the limitations of our use of it. To do this, we shall take reductionism itself to pieces to see how it works. For this, we need to use mathematical ideas: no cause for panic as there will be no equations and we shall take things quite slowly.

Reductionism is abstraction

Reductionism consists of the selection of a domain and the construction of simple ways of thinking about the domain. To a mathematician, both of these are forms of *abstraction*: think of it as simplification. Instead of thinking about the real world directly, you abstract from it to form a simplified model that you can reason about, and if you are really lucky, perform some calculations on.

To give you some feel for scientific abstractions, here are some examples.

1. The kinetic theory of gases as taught in schools is based on an abstraction of a gas as a collection of hard, perfectly elastic spheres, randomly moving about at high speed, whose collisions with the walls of the enclosing chamber account for the pressure the gas exerts on it. From a simple calculation of the mechanics of these collisions it is possible to derive a relation between the pressure, volume and temperature of a gas, which is experimentally observed as the gas laws. The domain here is any gas and the abstraction is to replace the molecules with rubber balls.
2. Simple school chemistry is based on a theory of matter as elements and compounds. Compounds are formed by combining elements, but the 92 naturally occurring elements cannot by chemical means be reduced to simpler forms. Thus iron and sulphur are elements. Iron is magnetic and sulphur burns. A mixture of iron filings and powdered sulphur can be separated with a magnet, but if you heat the mixture it forms ferrous sulphide a compound which is not magnetic and does not burn, but is useful for producing stink bombs. The 92 elements form the building blocks for all forms of matter and this is the basic idea behind chemistry. Chemistry is concerned with reactions between materials in the world: burning, rusting, dissolving, and all the other transformations that happen

when different substances are mixed together. For school chemistry, the simplification is to consider only the rules by which the elements combine without worrying too much about the underlying structure of matter which gives rise to the rules.

3. Biology deals with structures which are more complex than either physics or chemistry and the abstractions it uses are more difficult to grasp. The big idea behind biology is the theory of evolution. The abstraction here is that we are dealing with groups of organisms rather than individuals, and the response of these groups to the environment averaged over time. In reality, there is variation within the characteristics of the individuals making up the group and in the events happening to them, but it is the response of the group as a whole which defines the characteristics of the species. The domain of biology is made up of the living things. For evolution, the simplification is to deal with groups of individuals forming a species.

These simple examples can start to give some idea of the enormous number and range of abstractions in scientific thinking. They range from atoms and their constituents to whole galaxies of stars and from events lasting the merest fraction of a second to those taking millions of years. Abstraction is an essential element of any thinking about reality. It applies to every day life just as much as science. The doctor does not see a person coming into the consulting room, but a bunch of symptoms on legs. Airline booking clerks do not see people coming up to their desks but passengers who are to be separated from their luggage and sent on their way with the appropriate boarding card. But of course, being human, half the interest in these jobs occurs when the abstractions break down and when people step outside our expectations of them.

The scientific examples also reflect on one aspect of the reductionism versus holism controversy. If the argument is simply that reductionists miss interactions between the things they are focussing on and what lies outside the domain then this was understood centuries ago. You must have a domain which is large enough to encompass the phenomena you are studying. We have the prime example of the theory of evolution which is reductionist thinking but with a large scale focus. It is true that people often read more into a scientific theory than is warranted by the abstractions employed and again the theory of evolution is a prime example of this. I would say that many of our characteristics as humans can be explained by the theory of evolution, but that is not the end of the matter: we are not *simply* the product of evolution. For any scientific theory, the question arises as to whether it is a complete explanation and what can be legitimately deduced from it.

How abstraction works

Any scientific equation deals with an abstract world of symbols which represent real world quantities such as pressure and volume. Non-mathematical theories, like the theory of evolution, also work with concepts like gene pools which are abstract representations of aspects of reality. It is just that in these cases it is hard to get quantitative relations between the abstract concepts so mathematics is not as useful as it is, for example, in the theory of gravity. But whether one can write down an equation or not, there is an abstract world representing some aspects of the real world and the essential question is to explore the relation between the two.

This is illustrated in this diagram. The bottom cloud represents the real world in all its messiness, while the top cloud represents our beautiful abstract model of it. As a result of an event we are interested in, the real world changes from R1 to R2. Our model of the event in the abstract world causes it to change from A1 to A2. The connection between the real and the abstract worlds is represented by the abstraction function ***A***. Think of this as a recipe for getting from the real world to the abstract world. Now the abstraction is correct if, applying the same recipe to any event in the real world, you get the same before and after states for the corresponding event in the abstract world. In this case, mathematicians say the diagram commutes, which means that you can follow the arrows either way round the square to end up in the same place.

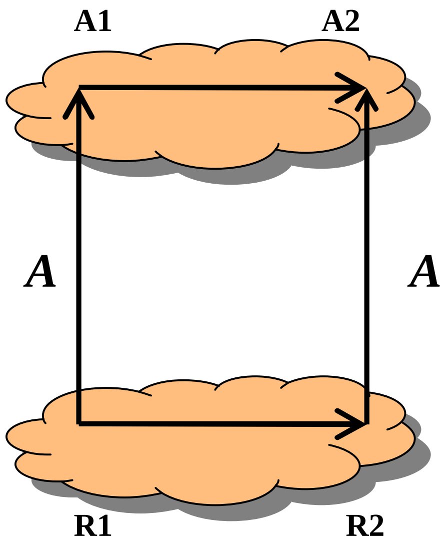

This diagram leaves mathematicians feeling pleased with themselves, but it is actually not much use to scientists. The problem is that one needs to test the theory and the test, that is the experiment, is carried out in the real world. The experiment needs to be carried out where the arrows meet, so that you can compare theory with experiment, so, as far as the scientists are concerned, the arrows meet in the wrong place. The experiment is carried out in the real world at R2, not A2 – that is just on paper. No problem, say the mathematicians, we can just turn the second arrow round to give this second commuting diagram.

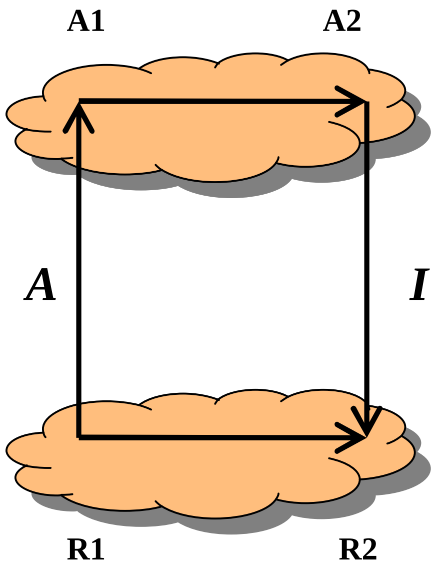

The second arrow ***I*** is now an interpretation function: it is a recipe for saying what real states correspond to what abstract states. ***I*** and ***A*** are related: it must be the case that if you apply the abstraction recipe to a real state and then immediately apply the interpretation, you should get back what you started with. Unfortunately, at this point the messiness of the real world strikes back: the abstraction is a simplification and the interpretation is the opposite. One abstract state corresponds to many different real world states. So the requirement is that if you interpret the abstraction the starting state should be one of the many possible ones the interpretation allows.

The number of states delivered by an interpretation can be, and usually is, astronomically large. If, for example, you are considering the properties of one hydrogen atom, the number of real states is the number of hydrogen atoms in the universe. There is no way of checking every single hydrogen atom. It is just an article of faith to scientists that an atom of hydrogen in the laboratory behaves in the same way as an atom on the other side of the universe. Astronomical observations based on this article of faith seem to give consistent results, so I would be very surprised if the faith were not justified, but for more complicated situations the interpretation of experimental results is always an issue.

The usual way of carrying out an experiment is to arrange it in such a way that the reality being tested is as close as possible to the abstract state. In physics, this usually involves extreme measures: reducing the temperature close to absolute zero; saturating the experiment with large electric and magnetic fields; isolating from vibrations; sending the experiment into space. There is surely not much original physics left which could be carried out in a school laboratory. But the use of these extreme measures does raise the questions of whether the result is an artefact of the experimental design and how significant the effect is in the real world.

In fields like biology and cosmology, the controlled experiment is a luxury, not often indulged in. It is not possible to replay a few million years of evolution or the moment of creation with a few different parameters to see what happens. Instead, you must rely on observations of the current state, to see how compatible they are with the abstract model you are dealing with. In other words, interpretation is everything and because of the nature of interpretation functions, you can only talk about probabilities. Science is never certain: it is just that scientific theories are as close to certainty as we are ever likely to get.

Medicine provides very good examples of the way in which science delivers answers which are not very certainly true, but usually very important. Medicine deals with a variety of abstractions. If it is a case of broken bones, doctors think in terms of articulated skeletons. In the case of contagious diseases, it is often important to think of the social aspects and the way in which people interact. For disorders like alcoholism on the other hand, the individual's spiritual situation needs to be considered: what are their hopes and fears and the pressures they undergo. But most people in a doctor's surgery are looking for a drug to cure their illness, rather than unwelcome news about a change in their lifestyle and it is interesting to see how this model of abstraction fits with the development of a new drug.

The mechanism for the operation of a drug usually involves the biochemistry of individual cells. Biochemists deal with the very large molecules which make up the structure of the organelles, the tiny components which constitute the living cell – itself of microscopic size. Biochemists view these macro-molecules rather like Lego building blocks which fit together to build up the cell. A drug operates by a chemical reaction with one of the macro-molecules, probably changing its shape and inhibiting some malfunction in the cell. Initially, one can test the operation of the drug on individual cells, possibly in a tissue culture, thereby taking into account interactions between cells. But the crucial test arises when the drug is tried out on people, so it is interesting to analyse this situation in terms of our model of reductionism.

In human trials, the real world is made up of the experimental subjects. This itself is an abstraction from the rest of the universe, but as we are not astrologers we shall leave the moon and planets out of it. The abstract world is made up of the cells being targeted in the drug trial and the event is the administration of the drug. There are two possible outcomes of the trial: either the drug had no effect, that is, R1 is the same as R2 or it had an effect, that is, R1 and R2 differ. Each of these cases can be further subdivided according to whether the change at the abstract level happened or not. If it did not, there are probably problems with the delivery of the drug or the susceptibility of the patients. If the drug actually reached the targeted cells and had an effect, but without changing the patient's medical condition, then the biochemical explanation of the illness is probably wrong. Further experiments might be necessary to distinguish these two cases.

Even where the drug does have an effect, it is still hard to draw conclusions. There are two cases again: the targeted cells did respond or they did not. If they did not, the response is induced by the experimental situation. If it was a good response, this is an example of the *placebo* effect. Some patients respond positively to any treatment given to them and experiments have to be designed to take this into account. Even if the targeted cells did respond, the change in outcome might still be due to the placebo effect, so it is still not sure that the biochemical explanation is a correct description of the illness.

The conclusion to draw is that scientific experiments do not deliver yes or no answers: at the best it is a question of probabilities. The other conclusion is that a scientific theory is only as good as the experiments which have verified it. A theory without experimental evidence is simply idle speculation, but even with a experimental data, a lot hangs on the quality of the experiment. Drugs trials involve very carefully controlled experiments, but even for these there can be fierce debate about the interpretation of the results. Other fields may not be so rigorous. Experiments testing evolutionary explanations of behaviour for example, are inherently difficult to carry out and are not subject to such rigorous review, so a measure of scepticism is needed when evaluating claims based on them.

This is not to say that science can tell us nothing in these areas. The link between smoking and lung cancer and global warming are typical cases involving uncertainty and difficult experiments. In these cases it is not possible to carry out laboratory experiments and one simply has to see how observations in the real world are compatible with the theories. Large numbers of observations are necessary and the evidence gradually mounts up, but it is always possible to dispute the conclusions. In both these cases there have been scientists prepared to dispute the theories, but as time has gone on, their number has dropped, usually ending up with small numbers with a vested interest. Nothing in this life is absolutely certain, but after a time, evidence from many experiments makes it a great deal easier to place your bets.

Science then may not be certain, but if the abstraction involved is used incorrectly it can be simply wrong, so we need to look at ways in which reductionism can fail.

What can go wrong with reductionism

An ill defined abstraction

People who play the lottery have a feeling that if certain numbers have come up frequently in the past, they are less likely to do so in the future and choose their numbers accordingly. This is a very human characteristic as we are always looking for patterns in events, in order to predict the future and do something about it. But it is not science. If you look at the machine which generates the numbers, you must have some idea as to how it can remember what has been selected from one choice to the next and the whole point of lottery machines is that they are carefully constructed so as not to do this. The reality of the machine does not lend itself to any model of its operation which could lead to one choice affecting another. Without that model, playing the numbers is superstition, not science.

Thought transference is another case where scientists are sceptical because of the qualities of the model. People communicate in all sorts of non-verbal ways, but not when confined within a sealed metal box. As a result of experiments of this nature one can say that the model of thought transference must use some force of nature which has not hitherto been discovered. This is rather a shaky model and as the results of experiments carried out to demonstrate thought transference are not conclusive, most scientists would say that thought transference does not exist. Some drugs can have similar marginal effects to what can be demonstrated in thought transference experiments: they usually do not make it to market, but if they have a model for their action, they are believable. Without it, they are much less likely to be investigated, still less used, and thought transference comes in to this category.

A different kind of criticism can be made of Richard Dawkins idea of a meme, modelled on a gene in evolutionary theory, to explain the spread and propagation of ideas. It is certainly obvious that ideas spread, but how exactly does a meme relate to reality? A gene is related to the genetic material carried from one generation to another, so what bit of reality does a meme correspond to? How and in what form is the meme expressed? Does it correspond to something like a scientific theory, in which case the meme might be a text book, or does it correspond to a fashion like wearing a baseball cap back to front, in which case it corresponds to images on television or in the newspapers. And what is the selection mechanism? With evolution the mechanisms are quite clear: individuals benefit from adaptive traits to the extent that they leave progeny. The ideas behind memes are so woolly, I would classify it as hand waving, not science.

Abstractions not containing the property being explained

Advances in science have made it possible to monitor the activity of the brain, even down to individual nerve impulses. These lead to some fascinating models of the brain based on interactions between individual nerve cells. As a result of studying injuries to the brain, it is, for example, possible to work our what parts of the brain are concerned with vision, with language, with the ability to do arithmetic, to recognise faces and a whole host of other activities which make up human life. This is leading people to think that here we can find an explanation for the elusive phenomenon of consciousness. Clearly, if you are looking at the firing of one nerve cell, you are not going to find anything about consciousness, any more than you could find out the properties of liquids by looking at an isolated atom. On the other hand, if you are looking at the firing of billions of neurons that make up the human brain, you are looking at an intractable problem. If you want to give an explanation of consciousness you need a higher level model and it is here where problems begin because this model must somehow contain the property being explored and we do not know exactly what that property is.

The problem can be illustrated with the history of the scientific development of the human concept of heat. In the eighteenth century, scientists were trying to define what heat was. Initially, it was thought that heat was a kind of fluid, “caloric”, which was contained within materials and which could be extracted in various ways. Experiments soon showed this to be wrong and eventually it was decided that heat was a form of energy, an energy consisting of the random motion of atoms, as opposed to other forms of energy such as those represented by a body in motion or the chemical energy stored in reactive materials. The interesting thing about heat is that it is easy to convert other forms of energy into heat, but difficult to convert heat into other forms of energy, leading to the very important science of thermodynamics.

Now we can accept this as an explanation for heat, because we can relate the scientific properties to our experience of heat: we can see how heat flows and why exercise makes us hot and why a breeze cools us down. But note in passing that physicists do not have the complete story on heat because to make these particular relations, you need to invoke physiology, chemistry and no doubt sociology too. One scientific explanation rarely provides a complete explanation of a human concept.

The problem with consciousness is that the phenomenon we are trying to explain is the way we perceive things, so what we are talking about is the way we perceive, the way we perceive.

Could you read that sentence again, please? In concrete terms, I have before me the image of my pencil making marks on paper. Without the firing of neurons in various parts of the brain, this image would not be there, but there is a very large gap between this image and the firing of neurons. One can make an analogy with a computer. The display on the computer screen mimics the pattern of electric charges in the video memory and at a higher level the charges within the computer's main memory, depending on what program is being run. But to talk about a computer in these terms may be adequate for a hardware engineer, but to miss out the user in front of the screen is really to miss the whole point of what the computer is about.

Brains are not computers, but our experience of consciousness is exactly like a little man inside our heads monitoring what is going on. This little man is, when it comes down to it, what we mean by "I". Now consciousness researchers have a habit of claiming that this little man does not exist, which means they are not looking at consciousness as most people experience it. And that experience is the one experience which is constantly before us in every waking moment of our lives. This is why consciousness research is hard: there are philosophical problems as well as the problem of explaining the most complicated object known to us. Simplifications which have thrown away "I" have thrown the baby out with the bath water and in this case we are not even sure what the baby is.

Theories without good evidence

It is a fundamental part of human nature to make theories about the world. Personally, I always take an umbrella when I particularly do not want it to rain on the theory that the weather demon, who is of course malevolent, will not know whether to send rain, and give me the satisfaction of having an umbrella, or not. Of course, I take my umbrella when it does look like rain anyway, so what this theory amounts to is that it often does not rain when it looks as if it is not going to rain and I happen to have my umbrella with me. Not good evidence for the theory by any stretch of the imagination.

This of course is not serious. But there are many theories about the world with about as much evidence to support them. Diets, for example. Here the problem is that your weight can change, but not as a result of the diet. Note that the model lying behind the idea of diets is perfectly sound. Your weight must be related somehow to what you eat, but the problem is with the exact composition and how that affects your weight. A good scientific experiment tries to ensure that the theory being tested is the only possible explanation for the change being observed and when it comes to tests on people that is difficult to achieve. A scientific theory is only as good as the evidence for it and one must always bear this in mind.

But this principle should also be applied to religion. If a religion is to say something about the real world it should be testable. Clearly such test would tend to disprove the idea of God as an old man on a cloud, receptive to requests from the faithful. But most religions these days involve aspects of our inner lives, so the tests we make are carried out by living. You have to live the religion, so faith is necessary to carry out the test. But those tests will involve values and goals and the idea of good, which I take to be aspects of reality.

The completeness of an abstraction

Is there anything more to say about the motion of heavenly bodies than what Newton and Einstein have said? Sitting in my study and thinking for all of five minutes, I would say not. But then doubts creep in. Some solutions of Einstein's general theory seem to allow time travel, which raises the question of whether causes precede effects and Hey! what does precede mean anyway? And what about quantum effects? Granularity seems to exist in the very fabric of space and the general theory does not take account of that. Each of these questions seems to raise doubts about relativity, which is one of the most experimentally verified theory in science. Relativity is not wrong, but that is not to say that it is a complete description of this aspect of reality. No model of reality tells the whole story and forgetting this fact is what most people mean when they use reductionism as a pejorative term.

And it is often deserved. Science is a very difficult intellectual discipline. You only make progress by concentrating on a small field and after a time, to some scientists that small field seems to be all there is – or at least, what lies outside is of no account, or merely detail, but certainly not a necessary component of our understanding of reality. But outside the realm of the professional scientist, reductionism in this bad sense is just as common. It is a good excuse: "It's not me, it's my genes", or my upbringing, or my biochemistry. One abstraction is never enough, it never tells the whole tale.

There is not test for the completeness of a theory only for its utility. If the theory answers the questions you are asking, it is good enough and that is how we make progress. Science after all, works. But a question usually leads to the requirement of another abstraction. For example, cholera is caused by a bacterium. End of story? No, if you want to stop an outbreak, you look at the mechanism of infection. The most famous case is the removal by John Snow of the handle of the Broad Street pump in Soho to contain an outbreak of cholera. That outbreak was cured by sociology, not biology. Later on, the civil engineering of London's sewers was even more effective.

Abstractions are approximations which only work in the domain they are designed for

Newton's laws of motion were revolutionary. A simple, intuitive and elegant idea expressed in new and beautiful mathematics explained the motion of the heavenly bodies. It is no wonder that this caused a sensation, not only among astronomers, but also the thinking public. Like the theory of evolution in the nineteenth century it raised the possibility that all of reality could now be understood. The impact of this thinking can be seen in Pope's *Epitaphs, To Sir Isaac Newton*:

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

And later, it can also be seen in William Blake's two paintings, *Sir Isaac Newton* and *The Ancient of Days*, where Blake is trying to reconcile his unorthodox religious views with this fundamental scientific insight.

For two centuries, Newton's laws of motion were the bedrock of science, so the impact of Einstein's theories of motion and gravity which challenged them can be imagined. Einstein's theories have been repeatedly verified to a high degree of accuracy, but we have not given up teaching and using Newton's laws. You only need to use relativity for very intense gravitational fields and very high speeds indeed. Newton's laws are an approximation. No doubt Einstein's are too. You need to know, when you are using the laws, what their limit of applicability is and this applies quite generally.

An interesting case arises with quantum theory, because of the way randomness enters in to the theory. Basically, the quantum equations of motion for fundamental particles only give the probability of a particle being detected at a given place or time. The entry of randomness at this fundamental level has almost as great an impact on the non-scientific world as Newton's laws themselves, but I believe the significance of this is misunderstood. Quantum theory applies to the domain of the exceedingly small. A quantum leap is the smallest change in energy it is possible to make, a fact which does not seem to have penetrated to those who use it as a figure of speech. In this microscopic domain, gravitational effects are simply ignored. In the macroscopic domain on the other hand, randomness is ignored because it can lead to no significant effects. I would say that anyone who thinks randomness provides a means whereby God can influence things is almost certainly barking up the wrong tree.

Evolution, as usual, presents other examples of the use of an abstraction outside its domain of application. The timescale for evolutionary change is very long. To be the product of adaptation, a characteristic ought to persist for thousands of years at least. Consequently you do not apply the theory of evolution to explain the different breeds of dogs and you should not apply the theory of evolution to human behaviour unless it is at a very fundamental level indeed.

Chance and necessity

Strangely enough, the most unreconstructed reductionists combine a belief in determinism according to scientific laws, with a belief in randomness, sometimes in effect creating gods of fortune and fate. If events are scientifically determined, where does the randomness come from?

The answer is that it arises from the abstraction. Focussing on one aspect of reality does not eliminate the rest of reality outside that focus, but the influence of it can often be represented by random noise, the effect of which can be taken into account by averaging. A horological friend of mine gave me the example of a pendulum, normally thought of as freely swinging in a uniform gravitational field. Other factors are thought of as noise, a source of error in clocks. In the search for more accurate time-keeping, one can broaden the abstraction to take into account temperature changes, often by some form of compensation. To reduce the noise further one can take into account the buoyancy and drag introduced by the air around the pendulum. What I find quite remarkable is that for a sufficiently accurate pendulum clock, the changes of rate can be analysed to show the influence of the tidal forces of the sun and moon on the gravitational field. No doubt tide charts could be provided for such a clock to correct it to show an even more accurate time, but there would still be noise in the rate arising from random vibrations from passing traffic and no doubt the occasional earthquake.

Abstraction is necessarily an approximation as some interactions have to be ignored And this bit of ignored reality always contains surprises, as the dinosaurs discovered to their cost when an asteroid collided with the earth at the end of the cretaceous era and the resulting climate change drove them to extinction. The theory of evolution does provides an interesting example of the way randomness can enter in. Evolution is a theory about the interaction of groups, that is, the individuals making up a species, with other groups and with the environment, averaged over considerable periods of time. The random event of one lion catching an antelope simply does not figure in the theory: it is what happens on average that counts. This average will be noisy in that the predation rate may show fluctuations from time to time, but these are sufficiently small not to invalidate the conclusions of the theory.

The asteroid strike on the other hand is a single event, not repeated. It would be wrong to draw the conclusion that we are adapted to asteroid strikes because the mammals survived the cretaceous-tertiary extinction. It was a “miracle”, that is, outside the scope of the theory. Now the interesting issue is what is the religious attitude to such events. In primitive religions, they were ascribed to the action of the gods, Jove sending thunderbolts, for example, but Christianity, and I suspect most modern religions, do not go along with this. Instead, what religion is concerned about is the attitude we should have in the face of these events. All things come to an end, and you should live your life knowing that. But if you escape something by the skin of your teeth do not sit back and congratulate yourself on your luck, but rather make sure it does not happen again.

There is a wonderful prayer of Reinhold Niebuhr which captures this:

> O God, grant us the serenity
> to accept what cannot be changed,
> the courage to change what can be changed,
> and the wisdom to know the difference.

Different religions balance these choices in different ways, but I would think that paganism and possibly atheism are much more likely to lead to fatalism.

Free will and determinism

Do we have free will? Newton started this question by revealing that the planets followed predetermined paths according to the laws of motion. Perhaps everything is determined in this way and the choices we make in life are predictable and free will is simply an illusion. Of course, if you accept this, almost the whole of human endeavour comes tumbling down including both science and religion, for if my discovery of a scientific law is simply the result of other scientific laws, how can you know that it is right? No one seriously uses this argument, but where exactly does it fall down? Exploring this paradox gives an insight into the interaction between different abstractions and helps us to understand where the properties of an abstraction in one field may constrain the type of model you can build in another. We shall start with the constraints you might acquire when one abstraction is described in mathematical terms. What exactly does follow from Newton's laws of motion or Schrödinger's quantum mechanical wave equation?

Mathematical determination

Non-mathematicians rarely come across equations, and the ones that do impinge on the public consciousness are very simple, like Einstein's equation $E = mc^2$, relating the energy E given out by the annihilation of mass m, with c being the velocity of light. Or at a more hum-drum level, there is Ohm's law, $V = IR$ relating the voltage V needed to drive a current I through a resistance R. These simple equations have the property that there is a single solution to them. Given m and the constant c, there is one value of E which satisfies the equation and only one current I which a voltage V can drive through a resistance R. However, most equations scientists deal with involve rates and here the solutions are not so simple. As an example consider a violin string. For small displacements of the string, the force which restores the string back to its resting position is proportional to the displacement of the string, so this instantly gives an equation which relates the acceleration of the string to the distance from its resting position. And this equation has a solution: the string passes its resting position at some speed, overshoots and then comes back again. In a word, it vibrates and at a frequency determined by the length of the string. This gives one solution of the equation, but there is also a solution at twice this frequency, three times it and so on. The actual motion of the string is governed by how it is bowed and stopped, so it is the human violin player who determines how the note sounds. The equations only determine the potential the string has to make a sound. Technically, we say that the actual solutions of mathematical equations involving rates of change are determined by what are called boundary conditions (in this case, the length of the stopped string and how it is bowed) and boundary conditions always flow from a higher level of abstraction than that treated by the equations.

Note in passing that when we are talking about causes, it is always with the idea of a human agent. A mechanical violin player would simply be moving a bow at a certain distance along the string and in this case we would say the note was determined by the designer of the automaton, or, if it was programmable, by the person (human!) feeding the program in. If you are talking about causes your abstraction must involve people.

Mathematics is not certain

It is one thing to write down an equation, but it is another matter entirely to solve it. And simple equations are surprisingly difficult to solve. The motion of three bodies under the laws of Newtonian attraction cannot be solved analytically and as for the solar system as a whole, don't even think of it. However, we manage to navigate rockets in very complicated trajectories, so there are obviously practical ways of achieving this, by approximation and correction. These methods achieved the stunning feat of navigating the Cassini space craft through Saturn's rings after a six year voyage, but they cannot answer questions like, is the solar system stable? We are fairly sure that we will not collide with Mars next July, or even within the next million years, but we cannot say anything about the next *billion* years other than that other factors, such as the evolution of the sun, will probably dominate matters by then.

Astronomical situations involve very long timescales, but uncertainty exists at the very smallest level. Schrödinger's equation, governing the quantum mechanical behaviour atomic particles, can be solved for only a few actual cases, most notably for the motion of a hydrogen atom, consisting of a proton and an electron only. But the tiniest droplet of ink, used in an inkjet printer, which is the smallest thing one could encounter in everyday life, contains millions of millions of molecules. No amount of computing power is going to solve the equations involved if you start at that level. Quantum mechanics is totally irrelevant to the engineering of an inkjet printer.

Both these equations have the advantage of being what mathematicians call linear: if you have two solutions of such an equation, the sum of the two is also a solution. This has a consequence that small changes of the variables only give rise to small changes in the solutions. For example, moving the earth's orbit outwards by a mile or two only makes the year a tiny bit longer, but will not make it travel in a figure of eight or leave the solar system entirely. Unfortunately, most equations which describe interesting properties of reality, like the weather, are non-linear and quite small changes in the parameters can lead to large changes in the solutions. This means the approximation and iteration methods used to tackle them have a much reduced limit of validity, which accounts for the inability to produce accurate weather forecasts for any time greater than a few days ahead.

Now it can be argued that although the equations that govern fundamental motions are quite incapable of solution, nevertheless, the equations are deterministic. (I shall ignore quantum indeterminacy which I believe to be a red herring.) This means that our actions are determined, even if unpredictable, and free will must be an illusion. Now I will argue that the equations themselves are approximations, as all abstractions with a limited domain of application are. Saying that a property of a low level abstraction applies at a higher level is a statement of faith, not a scientific fact. Taking a scientific law outside its domain of abstraction is wrong. To say that if you step over a hole you will certainly fall, follows from Newton's laws. But to say that if you could have measured all of Mozart's brain impulses you could predict he was about to compose the fortieth symphony is to talk nonsense.

At the risk of being repetitive, here are some further examples. Crick and Watson's spiral model of the DNA molecule was arrived at with stick and ball models of the atoms making up the molecule. The properties of the chemical bonds between the components of the molecule were understood intuitively, although they could be justified, in simpler situations, quantum mechanically. However, the spiral arrangement was justified on the basis of X-ray diffraction experiments. If someone had been able to do the quantum mechanical calculation and had come up with some structure which was not a spiral, every one would have assumed the calculation was wrong.

Another example from the same field was that genes were discovered independently, well before chromosomes and DNA were. There was certainly an "Aha!" moment when it was realised that the genes were embodied within stretches of the DNA molecule, but nothing in the laws of inheritance is dependent on the precise structure of DNA: the spiral structure which is important for cell division and growth, is largely irrelevant to the recombination of chromosomes which gives rise to genetic variation. Biochemists do not often solve Schrödinger's equation, but do often flip between different abstractions at a higher level than this, one of simplified molecular models and the other of genetic effects expressed in the working of cells.

It is essential to realise that these models are of equal status. There is an interaction between models when the levels of abstraction are very close. But when they are far apart, they are completely independent. Nothing in quantum mechanics is going to disprove the theory of evolution. And who pays any attention at all to quantum chromodynamics, describing the behaviour of elementary particles, apart from high energy physicists?

What free will is about

If the domain of an abstraction is important, we need to work out what is the domain to which free will applies. When is it reasonable to use this concept and when is it irrelevant? To explore this a little bit, consider the political sound bite "Tough on crime, tough on the causes of crime!" Now replace "crime" with "poverty": it does not make sense any more. The point is, people do not choose to be poor, but they do choose to commit a crime, sometimes, it is true, as a result of poverty. Free will is only an issue for moral choices: in fact, without free will, there can be no morals. It is wrong to *blame* the wolf for attacking the lamb, even if you take action to stop it.

We do, of course, make choices all the time, not just moral ones. This morning, for example, I had to choose whether to have marmalade or honey on my toast for breakfast. Free will in this case seems to be a bit of a heavy weight concept because no moral issues are involved. Sometimes the term free agency is used for those actions which are unpredictable, as opposed to free will which can be manifested even in a predictable decision. If you know someone has integrity you expect them to make just decisions, but when they do so, that does not mean they have not exercised their free will. So free will, as the term is usually applied, is about the moral choices we make, predictable or not.

So do we have free will?

We can summarise the argument so far along the following lines.

1. Low level, deterministic, models of reality only apply to very simple systems. Even where they do apply, actual behaviour is constrained by boundary conditions which are themselves determined by higher level abstractions

2. Higher level theories are verified independently of lower level ones, not derived from them. Where the levels of abstraction are close together, they can influence one another, but theories are only as good as their experimental validation, not their mathematical derivation. Thus when chemists study reactions at the molecular level, they are aware of the quantum mechanics of atoms, but the nature of the chemical bond is a separate theory, independently verified.

3. Abstractions are designed to answer questions and work within a particular domain. Abstractions are independent not only in the level of reality they are dealing with, but also in the concepts used. A concept appearing in one abstraction may not be present in another.

When we are talking about free will, we dealing with a bit of reality within ourselves and the only real issue of importance is for moral choices. It seems to be perverse to deny the existence of these moral choices, when, for most people, they are very keenly felt. They may be subjective, but simply saying choice is an illusion without further evidence is not science. So we need some model to explain the nature of free will.

Now I think it would be very difficult to argue that atomic theory, or nerve impulses, have anything to say about this aspect of reality: the abstractions which might have a bearing on the issue are much higher level than this and in this case one can only think that behavioural models, either genetically or environmentally based, are relevant. Clearly, the environment, in the form of the society and the upbringing of individuals is much the most important, so the question arises as to whether these environmentally based models could provide a *complete* explanation of the choices we make. And I would say not: people with similar upbringing do not make similar choices and behaviour generally is not consistent. The same person can respond differently on different occasions. Now is this a case of free agency or free will?

If you are looking at the choices people make from an environmental point of view, then choices made which are not constrained by the model will appear as noise, the operation of chance, and that is a product of this environmental abstraction. Do we need another abstraction to fully understand the problem? It is a religious point of view that we do, namely that moral choices are significant in themselves and can be judged relative to absolute standards of good and evil. One can see the effect by looking again at our sound bite. One can accept that "hard on the causes of crime" means looking at crime from an environmental point of view and doing something about the social conditions that give rise to crime. But "hard on crime" raises issues about whether one is simply working on deterrence (the environmental model) or judgement and repentance (the religious model). I believe that most people would understand and accept the latter model, without necessarily rejecting the former.

Is religion an abstraction?

If religion is an abstraction, and hence like science, it must be a way of thinking about reality. It need not be. A religion could be an exercise you practise to make you happy, to make you supple or to make you serene. I will leave it to the reader to decide what religions might fit these categories, but if you listen to the justifications people have for their religion it is often only in terms of their emotional states and the only reality such a religion can reflect is that state. A religion which you follow to make you feel happy can have nothing to say, for example, about the treatment of animals, the destruction of the environment or social justice.

Throughout history, most religions have been about aspects of reality which are much wider than this, the most obvious being fertility and sickness. Like all religions they should be judged not for the model of reality used as for how effective they are in their recipes for action in the light of this reality. A belief in evil spirits as a cause of sickness could have some validity if it gives rise to the right course of action. One would suppose that a religion which has been around for centuries must have some validity or it would not have persisted. However, in the case of fertility and sickness, scientific explanations are much more effective. We do not make human sacrifices to ensure fertility because fertilisers are much more effective. Moreover, in the light of other religions, for example Judaism, which deals with the reality of moral values, human sacrifice is wrong no matter how effective it was.

Religions today are about values and goals: what is good, what is evil, what is the purpose of life and how we should live. If this is an abstraction of reality, what is the nature of that reality? In particular, are these simply human properties like emotional states, or do values and goals reflect a reality which is outside human experience?

There is no doubt that religion is a human construction. Stones, plants and animals do not do religion. Animal actions are not good or evil, although one might constrain their actions for human purposes. But just because something is a human construction does not mean it has nothing to do with an external reality. Quite the reverse in fact – most human constructions do. Science and mathematics, for example, are largely to do with the physical world external to ourselves.

Certainly religions are worth studying from the human point of view. A quasi-evolutionary analysis, for example, could look at how religions grow and persist and the benefits they confer. Some people have even looked at nerve impulses hoping to find patterns there. But the issue is not with whether these abstractions are valid or not, but with whether they are complete. Does the fact that certain brain impulses occur when you understand Pythagoras's theorem have anything to do with whether the theorem is true or not? Does the fact that a religion might confer benefits in those who observe it mean that there is nothing more to say? Particularly as following some religions is onerous.

There is an interesting philosophical problem with any explanation of religions in human terms because there is then the problem of what to do with the explanation. If the only reason you have for following a religion, whose observance might be costly, is the benefit to future followers, then this is not a particularly compelling argument for incurring present cost. Any argument which does not make religion relate to an external reality will inevitably lead to people following their own desires and interests, leading to a world of slobs controlled by people in power. Most people would say that is not a desirable outcome and misses aspects of reality as they see it.

Thinking of religion as an abstraction in this way shows where the difficulties lie. When religions deal with external matters such as sickness or fertility, the tests you apply can be carried out objectively. But when religions are concerned with values and goals the tests can only be carried out subjectively. Values and goals are concerned with the choices you make in life. You need commitment to make difficult choices and you need commitment to decide what you are doing with your life. I would say that everyone has a religion because they accept some values and goals. People change their religion when they find they can no longer live according to them.

An abstraction is irrelevant unless it can be tested against reality. A religion of values and goals is about that bit of reality concerned with how we live. As a result a religion has to be applicable to the lives of all sorts of people and the models they employ have to be simple enough to be universally understood. A religion which was as subtle an abstraction as the theory of relativity would be useless as most people would simply not understand it. Because of this simplification, it is particularly important not to take religious beliefs outside their domain of application, which in this case is concerned with how we live. When we talk about God, it is idle speculation because you are not going to test it unless it leads to some choice for life. Realising this point will resolve many paradoxes. Do not think what "God" means from God's point of view (that is, by thinking about what God is feeling or how God acts) but rather from what consequences there are for our behaviour. Two people of different religions who end up with the same behaviour are in more agreement about the nature of God than two people of the same religion who behave differently.

The big issue for the religion as abstraction debate is whether this idea of God is a characteristic of our human nature or whether the idea does relate to some other aspect of reality. I would say it does, but this can only be a statement of faith, because you have to have faith to try it out.