What is missing from Popper's evolutionary epistemology is an argument that links the fitness of a theory to its truth. Popper's broader philosophical views offer scant hope of helping us formulate such an argument. He claims that when a conjecture is at odds with our data, we can reject it. Under these circumstances the conjecture has been 'falsified'. But he denies that we have good reason to believe a conjecture that has not been falsified is true. Indeed, he rejects any extrapolation from a theory's past successes in avoiding falsification to its prospects for future success against the tribunal of experience. Popper's picture of science is made more complex by his further claim that statements of the data, as well as theoretical hypotheses, have the status of conjectures. Since all conjectures are (for Popper) wholly tentative, if a theory avoids falsification this means only that one conjecture – the hypothesis – is consistent with another set of tentative conjectures – the data. For Popper, then, science is a process by which one set of tentative conjectures becomes adapted to another set of tentative conjectures. He offers no convincing arguments for why such a state of adaptedness has any bearing on the truth of either set of statements.

5. MEMES

In recent years, theorists from different disciplines have proposed evolutionary models of the sciences which take the loose analogy Popper draws between science and natural selection, and bring it into far closer formal alignment with the principles of modern evolutionary biology. Such models have their roots in works by the psychologist D. T. Campbell, who regards the growth of knowledge as a process of what he calls 'blind-variation-with-selective-retention' (Campbell 1974). They also owe a lot to Richard Dawkins' speculative remarks about the possibility of non-genetic evolution at the end of *The Selfish Gene* (1976), and to the philosopher of biology David Hull's pioneering studies of scientific change (Hull 1988).

In order to examine these more formal models of scientific evolution, we need to say a little about Hull's distinction between *replicators* and *interactors* (ibid.: 408). This is closely related to Dawkins' distinction between *replicators* and *vehicles*. When we introduced natural selection back in chapter two, we noted the widely acknowledged definition of selection as a process operating on entities that vary in their fitness, and which reproduce in such a way that offspring resemble parents. These conditions are stated in an abstract way, allowing that any set of entities, whether they are organisms, computer viruses, ideas or artworks, might be said to undergo selection, just so long as they reproduce, and offspring resemble parents. The virus in my computer is the 'offspring' of the virus in the computer that was the source of the infection; the scepticism I have about the existence of God is the 'offspring' of the atheism of David Hume, whose *Dialogues Concerning Natural Religion* converted me; the picture I painted is the 'offspring' of the Monet from which it was copied.

Hull claims that all selection processes – standard organic evolution included – require entities that play two distinct roles, roughly corresponding to the twin requirements that offspring resemble parents, and that parents differ in their fecundity. *Replicators* are entities that copy themselves, thereby ensuring trans-generational resemblance. Genes are usually thought of as replicators in organic evolution: offspring resemble parents, so it is said, because genes have the ability to make copies of themselves. *Interactors* are entities which cause replicators to appear in different proportions in the offspring generation, in virtue of their interactions with the environment. Fast-running wolves catch deer more efficiently, and as a result of this their genes are copied in greater proportions than the genes of slow-running wolves. In this particular case, wolves are interactors, while wolf genes are replicators; Hull's view allows that under some circumstances, a single type of entity (an asexually reproducing bacterium, for example) might act as replicator and interactor at the same time.

The replicator/interactor distinction raises many interesting questions that will not be addressed here. How, precisely, should we define these terms (Griffiths and Gray 1994)? Are genes the only replicators in organic evolution (Sterelny *et al.* 1996)? Could selection occur without replicators (Godfrey-Smith 2000)? If not, how can natural selection explain the initial emergence of replicators, which are clearly complex entities in their own right? Rather than looking at these questions in detail, let us instead look at Hull's application of the replicator/interactor distinction to scientific evolution. He claims that stability over time in science – whether that is stability of the theories believed by successive generations of scientists, or of styles of production of scientific documents, or of techniques used in the lab – should be explained by citing the transmission of reliably copied replicators. But the replicators in question are not genes: 'The mode of transmission in science is not genetic but cultural, most crucially linguistic. The things whose changes in relative frequency constitute conceptual change in science as elsewhere are "memes", not genes . . . ' (Hull 2001: 98).

Memes are cultural replicators. Richard Dawkins, who coined the term, gives us a list of exemplary memes which includes: 'tunes, ideas, catch-phrases, clothes fashions, ways of making pots or of building arches' (Dawkins 1976: 206). Let us focus for the moment on tunes. These are indeed particularly 'catchy' bits of culture, which often hop from person to person. A friend hums *Singin' in the Rain* one morning. I hear the tune, and find myself whistling it later on. By the evening, five or six of my colleagues have the same tune in their heads and at their lips. Thus the replicator makes copies of itself, and various interactions between noises and human brains cause this spread. The tune of *Singin' in the Rain*, Dawkins says, is a meme. Scientific theories, Hull says, are memes, too. They also make copies of themselves, hopping from brain to brain. Their rates of spread depend on the effects of diverse interactors – books, articles, conversations, tools – on the scientific environment. The criteria, conscious or unconscious, which scientists themselves bring to bear when assessing the merits of a theory constitute an important set of features of the selective environment of memes.

The meme concept has attracted a fair amount of hostile commentary. Even supporters of the evolutionary view of culture often argue that formalising cultural evolution using memes is misguided. Let me briefly review three fairly frequent complaints directed at the meme/gene analogy, before going on to pinpoint what is the most serious problem affecting memetics.

'Genetic units are discrete particles; culture is not composed of discrete units'

Critics sometimes argue that it makes no sense to think of an idea as a gene-like unit, which can be analysed in isolation. Ideas come packaged as interconnected systems – the idea of a god, for example, can only be understood when one also understands other ideas to which it is related. Depending on which religion we are discussing these might be ideas relating to paternity, maternity, grace, knowledge, love, vengefulness, and so forth. This also means that one cannot treat all instances of belief in a god as instances of the same type of meme. If belief in God only makes sense in the context of the system it features in, then one religion's 'belief in God' is only superficially similar to that of another. The anthropologist Adam Kuper summarises: 'Unlike genes, cultural traits are not particulate. An idea about God cannot be separated from other ideas with which it is indissolubly linked in a particular religion' (Kuper 2000: 180).

The memeticist is likely to accept that ideas need to be understood in context, and that not every belief in god is the same type of belief. But she is likely to add that genes are just like ideas in these respects. She might point out that genes depend for their effects on interactions with other genes in the organism. She might add that superficially similar genes, identified by their DNA sequences, can have very different roles in different species, so that it makes little sense to think of them as instances of the same type for the purpose of evolutionary analysis. Genes do not have a life of their own, in isolation from their specific web of relations to other genes, any more than ideas do. Even so, once we have specified some particular species, perhaps even some particular population, one can isolate the role of a gene in that context. The same, she might say, goes for memes. Ideas can be assigned individual roles once we specify a particular social context for investigation.

'Genetic units makes copies of themselves; cultural units do not'

Memes are supposed to be replicators. They are supposed to make copies of themselves. Now it is certainly true that ideas spread through groups of people. But it is less clear whether they do so by making copies of themselves. The anthropologist Dan Sperber complains that:

> . . . most cultural items are 're-produced' in the sense that they are produced again and again—with, of course, a causal link between all these productions—but are not reproduced in the sense of being copied from one another . . . Hence they are not memes, even when they are close 'copies' of one another (in a loose sense of 'copy', of course).
>
> (Sperber 2000: 164–65)

Recall the example of my whistling *Singin' in the Rain*. I whistle it because I heard the tune earlier in the morning. In a sense, a reproduction has been made of the tune. But although my performance resembles the earlier one, is mine copied from the earlier one? Perhaps it is: perhaps I listen very carefully to the tune, take efforts to memorise it, and whistle it myself. But probably I do not do this. More likely I hear a little of the tune and think 'Aha! That chap is humming *Singin in the Rain*! Such a fine tune!' The tune is already familiar to me, I have no need to listen carefully, and I begin to whistle it myself. In this second case it is somewhat strained to say that my version of the tune is a *copy* of the one I hear. It seems more appropriate to say that my hearing the tune *triggers* the performance of a tune that is already in my repertoire. Sperber understands replication as copying in the strong sense, rather than as the triggering of a resembling performance. He goes on to argue that most cultural reproduction is of the triggering type, not the copying type. As a result replication is comparatively rare, and there are fewer instance of meme-like reproduction than at first meet the eye. Spelling out in precise terms what the difference is between copying and triggering may be difficult, but Sperber is right to remind us that there are many different ways for the same idea, or the same behaviour, or the same tune, to be 'reproduced' through a population. What is more, if the successes of organic evolutionary theory are anything to go by, evolutionary theory becomes enlightening when we are able to characterise in some detail the modes that reproduction takes. This is why Mendel's genetic laws are important: Mendel's laws tell us something about the general patterns of parent/offspring relations, which in turn help us to explain the makeup of successive generations of a population. Our theories of cultural evolution, if they are to be enlightening, need to do more than assert that culture contains varied ideas which are reproduced at different rates. In this sense, the mere claim that culture evolves is not sufficient to make cultural evolutionary theory informative. These theories also need a rich enough vocabulary to capture different modes of cultural reproduction, and they need to investigate how those modes of cultural reproduction affect the composition of successive cultural generations.

'Genes form lineages; cultural units do not'

A worry that is closely related to Sperber's draws on the fact that ideas do indeed spread through populations when individuals learn from each other, but these ideas do not always form *lineages* in the ways that genes do (Boyd and Richerson 2000). In principle, I could look into my genome and say (for most of my genes, at least), which came from my father and which from my mother. Each gene is derived from a single individual, in such a way that we might trace a lineage back through time. Can we do this for cultural items? Not always. Consider my knowledge of the tune of *Singin' in the Rain*. It is unlikely that there is a single source from which this knowledge is derived. It is unlikely, for example, that I learned this tune because one other person whistled it. I probably picked the tune up over time, from exposure to parents, friends, various showings of films, and so forth. The knowledge of a tune like *Singin' in the Rain* spreads through a population, and various facts might make it spread more quickly than other tunes. Yet it is misleading to call *Singin' in the Rain* a meme, because unlike genes, people who know the tune of *Singin' in the Rain* have rarely inherited it from a single individual. What facts might make one tune more likely to spread than others? In part, of course, we can point to facts that make an individual who knows the tune more likely to whistle it, and to facts that make an individual who hears it more likely to remember it. But a tune could score comparatively poorly on these characteristics and still spread faster than its competitors simply because it is ubiquitous. If a record company ensures that a melody is played through all available radio and TV networks, then even a tune that is comparatively un-catchy will quickly become known by millions. This underlines an important limitation for memetics. In organic evolution, the swift spread of some variation through a population typically indicates that the variation in question confers high reproductive success on its bearers. Things are more complicated at the cultural level. We cannot infer from the swift spread of a tune through a population that the tune has features than make it likely to hop from mind to mind. The tune may not be especially 'contagious' or 'catchy' at all; the tune's producers may just be powerful enough to make it ubiquitous, hence more likely to be learned than far catchier, but more poorly-funded, competitors. Once again, it is important that our cultural evolutionary theories are rich enough to document the diverse reasons why an idea may spread, and the memetic theory, by drawing a very close analogy between organic and cultural evolution, threatens to obscure the important distinction between contagious and power-assisted spread.

6. CULTURAL EVOLUTION WITHOUT MEMES

Martin Gardner (quoted in Aunger 2000: 2) complains that: 'memetics is no more than a cumbersome terminology for saying what everybody knows and that can be more usefully said in the dull terminology of information transfer'. Even if memeticists are right that culture evolves, this is not informative unless they go some way to spelling out the details of how culture evolves. The memeticist might answer with a case-study. Suppose we are trying to explain why more people buy new Minis than buy new Beetles. We could do so by suggesting that one meme – the inclination to buy a new Mini – is fitter than another – the inclination to buy a new Beetle. What makes one fitter than the other? Perhaps Minis look cooler than Beetles, or perhaps they run better, or they are cheaper. The result is that more copies are made of the Mini-purchasing meme than the Beetle-purchasing meme, and these memes cause differential purchasing behaviour on the part of their bearers. If case-studies like this are the best we can come up with, Gardner's objection is reinforced. This memetic explanation is merely a cosmetic repackaging of the kind of story an economist, or a psychologist, might tell about why more people buy Minis than Beetles.

The assertion that cultural change can be understood in terms of various factors that explain the relative successes and failures of different ideas verges on the trivial. A useful theory of cultural evolution needs to offer some insight regarding what factors need to be taken into account in explaining the changing composition of a population of ideas, and these factors need to be unlikely to be noticed by students of more traditional disciplines like economics, or psychology. Memetics is particularly unlikely to yield an informative cultural evolutionary theory of this kind. Its proponents appear to think that because genetic models of evolution have been largely successful in the organic realm, similar models must be the best ones to use for the cultural realm. Genes are understood as discrete particles that are faithfully copied; consequently memes are understood as discrete particles that are faithfully copied. Perhaps memeticists think that the pioneers of the modern theory of evolution – people like R. A. Fisher – showed that natural selection can only work when inheritance is 'particulate'. Remember the problem we examined in chapter two, which Fleeming Jenkin raised for natural selection. If offspring are always intermediate in character between their parents, then, Jenkin said, it seems that beneficial mutations will not be added up and preserved by selection, but instead they will be washed away over time by the action of 'blending inheritance'. Fisher did not solve this problem by arguing that natural selection can work only if inheritance is underpinned by the transmission of discrete particles (as opposed to a general blending of parental traits). Instead, Fisher argued that natural selection would only be able to produce cumulative adaptation within a system of 'blending' inheritance if mutation rates were very high – certainly higher than observed genetic mutation rates. Fisher's achievement was to show how a system of particulate inheritance would enable natural selection to operate effectively even with low genetic mutation rates.

Fisher's work immediately prompts a series of questions we can ask about cultural evolution. What might we mean by the 'rate of cultural mutation'? How could we measure it? Is the cultural mutation rate higher than the genetic mutation rate? Is it high enough for natural selection to be able to operate effectively without the faithful replication of cultural 'particles'? Is each cultural trait of an individual organism in fact a blend, assembled from diverse influences (such as parents, siblings and authority figures)? How does cultural inheritance affect the natural selection of organisms? How, for example, does cultural inheritance – which can perhaps maintain the presence over time of traits in social groups – affect the operation of natural selection at the level of the group? All of these questions are best answered using a combination of rigorous statistical modelling coupled to detailed empirical investigation – just the techniques that enabled the pioneers of the modern synthesis to make natural selection a well-understood and well-confirmed explanation for organic evolution. The memetic view has a tendency to obscure the importance of questions like these. That is why the most constructive work in cultural evolutionary theory has been done by those who are sceptical of memes – people like evolutionary anthropologists Robert Boyd and Peter Richerson, who try to answer just these questions (Richerson and Boyd 2005).

To give a hint of the promise that Boyd and Richerson's approach holds for the understanding of culture, consider their discussion of technological innovation (ibid.: 52–54). They begin by telling the story of the development of the modern ship's compass. It is a complex one, which starts with the discovery that naturally-occurring magnetite has a tendency to directional orientation.

Further refinements are spread over centuries and continents. They include the production of magnetite needles that can be floated in a bowl of water, the mounting of a magnetic needle on a vertical pin bearing, the addition of iron balls that cancel out the distorting effects of steel-hulled ships, and the perfection of various systems that damp the response of the compass to the ship's rocking.

This account, while recognisably evolutionary in its commitment to a Darwinian form of gradualism, is unlikely to startle any historian of technology. It is no surprise to learn that innovation often proceeds through the accumulation of many small steps, which have taken place over a considerable breadth of time and space. Things get more interesting when Boyd and Richerson begin to ask comparative questions. European empire-builders successfully invaded the Americas; the Americans did not invade Europe. Why did it happen this way round? They follow Jared Diamond (1996) in attributing it to the greater pace of technological innovation in European, compared with American, societies. In chapter two we noted Darwin's own recognition – a purely statistical insight – that the size of a social group can affect the chances of useful technological inventions being produced in it. Boyd and Richerson offer a similar form of explanation for the difference in innovative pace in the two landmasses, which draws once again on Diamond's work. They consider innovation a function of the likely rate at which cultural mutation can be generated, and of the chances of advantageous cultural mutations increasing in frequency once produced. They suggest that:

> . . . the greater size of the Eurasian continent, coupled with its east-west orientation, meant that it had more total innovation per unit of time than smaller land masses, and that these innovations could easily spread through the long east-west bands of ecologically similar territory.
> (Richerson and Boyd 2005: 54)

Suggestions like this are certainly speculative, but they show the potential for explanatory novelty offered by the evolutionary approach to culture.

SUMMARY

Darwin's work has influenced epistemology in two main ways. The first is direct. Prior to Darwin, philosophers had long been divided on the existence of innate knowledge. On the one hand, it seemed to many that we know things that we have not had to learn. But if learning does not account for the possession of a true belief, what does? Darwin's evolutionary theory immediately suggests a plausible mechanism that might explain this – namely, natural selection – and his work thereby offers the promise of rendering innate knowledge respectable. Darwin's work has also influenced epistemology in an indirect fashion. Natural selection can be stated in an abstract way, which allows us to see entities other than organisms as subject to selection processes. Entities of any type can be said to evolve by natural selection, so long as they vary, they reproduce and offspring resemble parents. So-called 'evolutionary epistemology' claims that scientific theories meet these conditions, and it consequently studies scientific change as an evolutionary process. Recently, evolutionary epistemologists have made widespread use of the meme concept, regarding scientific theories, and ideas in general, as memes. Memes are supposed to be cultural analogues of genes. They are replicators – that is, entities which make copies of themselves – and they underlie cultural inheritance. We saw in this chapter that there are reasons to be sceptical of the meme/gene analogy. It is far from clear that all ideas are replicators (although some might be), and it also unlikely that ideas always form lineages (although sometimes they do). More importantly, even if memetics' defenders are right to say that culture evolves, and that cultural evolution consists in the differential spread of different types of meme, it is unclear how much insight this brings that could not be had just as well by using models from psychology or economics. This is not to say that no cultural evolutionary theory has value, but such theories need to examine how ideas are reproduced, how they mutate, how the structure of a population of ideas affects the prospects of that population, and so forth. These are the kinds of questions that needed to be answered before Darwin's theory of natural selection could be applied in a detailed manner to organic evolution, and the same questions need to be asked in the cultural realm.

FURTHER READING

Elliott Sober has a clear discussion of the issues surrounding the relative evolutionary merits of belief-forming mechanisms liable to error and those that are more accurate, as well as the merits of innate belief compared with learning:

Sober, E. (1994c) 'The Adaptive Advantages of Learning and A Priori Prejudice', in E. Sober, *From a Biological Point of View*, Cambridge: Cambridge University Press.

A good place to turn to for Popper's evolutionary epistemology is:

Popper, K. (1962) 'Conjectures and Refutations' in K. Popper, *Conjectures and Refutations*, London: Routledge.

An extended statement of Campbell's evolutionary epistemology appears in:

Campbell, D. T. (1974) 'Evolutionary Epistemology', in P. Schilpp (ed.) *The Philosophy of Karl Popper*, LaSalle, IL: Open Court.

A useful collection of essays on memes – some hostile, some friendly – was published a few years ago:

Aunger, R. (2000) *Darwinizing Culture*, Oxford: Oxford University Press.

Daniel Dennett also discusses the meme concept in a balanced and characteristically lively fashion towards the end of his book on Darwinism:

Dennett, D. (1995) *Darwin's Dangerous Idea: Evolution and the Meanings of Life*, London: Allen Lane.

A particularly sophisticated overview of modern cultural evolutionary theory can be found in:

Richerson, P. and Boyd, R. (2005) *Not by Genes Alone: How Culture Transformed Human Evolution*, Chicago: University of Chicago Press.