

Mememes: the new replicators

So far, I have not talked much about man in particular, though I have not deliberately excluded him either. Part of the reason I have used the term 'survival machine' is that 'animal' would have left out plants and, in some people's minds, humans. The arguments I have put forward should, *prima facie*, apply to any evolved being. If a species is to be excepted, it must be for good particular reasons. Are there any good reasons for supposing our own species to be unique? I believe the answer is yes.

Most of what is unusual about man can be summed up in one word" 'culture'. I use the word not in its snobbish sense, but as a scientist uses it. Cultural transmission is analogous to genetic transmission in that, although basically conservative, it can give rise to a form of evolution. Geoffrey Chaucer could not hold a conversation with a modern Englishman, even though they are linked to each other by an unbroken chain of some twenty generations of Englishmen, each of whom could speak to his immediate neighbours in the chain as a son speaks to his father. Language seems to 'evolve' by non-genetic means, and at a rate which is orders of magnitude faster than genetic evolution.

Cultural transmission is not unique to man. The best non-human example that I know has recently been described by P. F. Jenkins in the song of a bird called the saddleback which lives on islands off New Zealand. On the island where he worked there was a total repertoire of about nine distinct songs. Any given male sang only one or a few of these songs. The males could be classified into dialect groups. For example, one group of eight males with neighbouring territories sang a particular song called the CC song. Other dialect groups sang different songs. Sometimes the members of a dialect group shared more than one distinct song. By comparing the songs of fathers and sons, Jenkins showed that song patterns were not inherited genetically. Each young male was likely to adopt songs from his territorial neighbours by imitation, in an analogous way to human language. During most of the time Jenkins was there, there was a fixed number of songs on the island, a kind of 'song pool' from which each young male drew his own small repertoire. But occasionally Jenkins was privileged to witness the 'invention' of a new song, which occurred by a mistake in the imitation of an old one. He writes:

'New song forms have been shown to arise variously by change of pitch of a note, repetition of a note, the elision of notes and the combination of parts of other existing songs ... The appearance of the new form was an abrupt event and the product was quite stable over a period of years. Further, in a number of cases the variant was transmitted accurately in its new form to younger recruits so that a recognizably coherent group of like singers developed.' Jenkins refers to the origins of new songs as 'cultural mutations'.

Song in the saddleback truly evolves by non-genetic means. There are other examples of cultural evolution in birds and monkeys, but these are just interesting oddities. It is our own species that really shows what cultural evolution can do. Language is only one example out of many. Fashions in dress and diet, ceremonies and customs, art and architecture, engineering and technology, all evolve in historical time in a way that looks like highly speeded up genetic evolution, but has really nothing to do with genetic evolution. As in genetic evolution though, the change may be progressive. There is a sense in which modern science is actually better than ancient science. Not only does our understanding of the universe change as the centuries go by: it improves. Admittedly the current burst of improvement dates back only to the Renaissance, which was preceded by a dismal period of stagnation, in which European scientific culture was frozen at the level achieved by the Greeks. But, as we saw in Chapter 5, genetic evolution too may proceed as a series of brief spurts between stable plateaus.

The analogy between cultural and genetic evolution has frequently been pointed out, sometimes in the context of quite unnecessary mystical overtones. The analogy between scientific progress and genetic evolution by natural selection has been illuminated especially by Sir Karl Popper. I want to go even further

into directions which are also being explored by, for example, the geneticist L. L. Cavalli-Sforza, the anthropologist F. T. Cloak, and the ethologist J. M. Cullen.

As an enthusiastic Darwinian, I have been dissatisfied with explanations that my fellow-enthusiasts have offered for human behaviour. They have tried to look for 'biological advantages' in various attributes of human civilization. For instance, tribal religion has been seen as a mechanism for solidifying group identity, valuable for a pack-hunting species whose individuals rely on cooperation to catch large and fast prey. Frequently the evolutionary preconception in terms of which such theories are framed is implicitly group-selectionist, but it is possible to rephrase the theories in terms of orthodox gene selection. Man may well have spent large portions of the last several million years living in small kin groups. Kin selection and selection in favour of reciprocal altruism may have acted on human genes to produce many of our basic psychological attributes and tendencies. These ideas are plausible as far as they go, but I find that they do not begin to square up to the formidable challenge of explaining culture, cultural evolution, and the immense differences between human cultures around the world, from the utter selfishness of the Ik of Uganda, as described by Colin Turnbull, to the gentle altruism of Margaret Mead's Arapesh. I think we have got to start again and go right back to first principles. The argument I shall advance, surprising as it may seem coming from the author of the earlier chapters, is that, for an understanding of the evolution of modern man, we must begin by throwing out the gene as the sole basis of our ideas on evolution. I am an enthusiastic Darwinian, but I think Darwinism is too big a theory to be confined to the narrow context of the gene. The gene will enter my thesis as an analog), nothing more.

What, after all, is so special about genes? The answer is that they are replicators. The laws of physics are supposed to be true all over the accessible universe. Are there any principles of biology that are likely to have similar universal validity? When astronauts voyage to distant planets and look for life, they can expect to find creatures too strange and unearthly for us to imagine. But is there anything that must be true of all life, wherever it is found, and whatever the basis of its chemistry? If forms of life exist whose chemistry is based on silicon rather than carbon, or ammonia rather than water, if creatures are discovered that boil to death at —100 degrees centigrade, if a form of life is found that is not based on chemistry at all but on electronic reverberating circuits, will there still be any general principle that is true of all life? Obviously I do not know but, if I had to bet, I would put my money on one fundamental principle. This is the law that all life evolves by the differential survival of replicating entities.* The gene, the DNA molecule, happens to be the replicating entity that prevails on our own planet. There may be others. If there are, provided certain other conditions are met, they will almost inevitably tend to become the basis for an evolutionary process.

But do we have to go to distant worlds to find other kinds of replicator and other, consequent, kinds of evolution? I think that a new kind of replicator has recently emerged on this very planet. It is staring us in the face. It is still in its infancy, still drifting clumsily about in its primeval soup, but already it is achieving evolutionary change at a rate that leaves the old gene panting far behind.

The new soup is the soup of human culture. We need a name for the new replicator, a noun that conveys the idea of a unit of cultural transmission, or a unit of *imitation*. 'Mimeme' comes from a suitable Greek root, but I want a monosyllable that sounds a bit like 'gene'. I hope my classicist friends will forgive me if I abbreviate mimeme to *meme** If it is any consolation, it could alternatively be thought of as being related to 'memory', or to the French word *meme*. It should be pronounced to rhyme with 'cream'.

Examples of memes are tunes, ideas, catch-phrases, clothes fashions, ways of making pots or of building arches. Just as genes propagate themselves in the gene pool by leaping from body to body via sperms or eggs, so memes propagate themselves in the meme pool by leaping from brain to brain via a process which, in the broad sense, can be called imitation. If a scientist hears, or reads about, a good idea, he passes it on to his colleagues and students. He mentions it in his articles and his lectures. If the idea catches on, it can be said to propagate itself, spreading from brain to brain. As my colleague N. K. Humphrey neatly summed up an earlier draft of this chapter: '... memes should be regarded as living

structures, not just metaphorically but technically.* When you plant a fertile meme in my mind you literally parasitize my brain, turning it into a vehicle for the meme's propagation in just the way that a virus may parasitize the genetic mechanism of a host cell. And this isn't just a way of talking—the meme for, say, "belief in life after death" is actually realized physically, millions of times over, as a structure in the nervous systems of individual men the world over.'

Consider the idea of God. We do not know how it arose in the meme pool. Probably it originated many times by independent 'mutation'. In any case, it is very old indeed. How does it replicate itself? By the spoken and written word, aided by great music and great art. Why does it have such high survival value? Remember that 'survival value' here does not mean value for a gene in a gene pool, but value for a meme in a meme pool. The question really means: What is it about the idea of a god that gives it its stability and penetrance in the cultural environment? The survival value of the god meme in the meme pool results from its great psychological appeal. It provides a superficially plausible answer to deep and troubling questions about existence. It suggests that injustices in this world may be rectified in the next. The 'everlasting arms' hold out a cushion against our own inadequacies which, like a doctor's placebo, is none the less effective for being imaginary. These are some of the reasons why the idea of God is copied so readily by successive generations of individual brains. God exists, if only in the form of a meme with high survival value, or infective power, in the environment provided by human culture.

Some of my colleagues have suggested to me that this account of the survival value of the god meme begs the question. In the last analysis they wish always to go back to 'biological advantage'. To them it is not good enough to say that the idea of a god has 'great psychological appeal'. They want to know *why* it has great psychological appeal. Psychological appeal means appeal to brains, and brains are shaped by natural selection of genes in gene-pools. They want to find some way in which having a brain like that improves gene survival.

I have a lot of sympathy with this attitude, and I do not doubt that there are genetic advantages in our having brains of the kind that we have. But nevertheless I think that these colleagues, if they look carefully at the fundamentals of their own assumptions, will find that they are begging just as many questions as I am. Fundamentally, the reason why it is good policy for us to try to explain biological phenomena in terms of gene advantage is that genes are replicators. As soon as the primeval soup provided conditions in which molecules could make copies of themselves, the replicators themselves took over. For more than three thousand million years, DNA has been the only replicator worth talking about in the world. But it does not necessarily hold these monopoly rights for all time. Whenever conditions arise in which a new kind of replicator *can* make copies of itself, the new replicators *will* tend to take over, and start a new kind of evolution of their own. Once this new evolution begins, it will in no necessary sense be subservient to the old. The old gene-selected evolution, by making brains, provided the soup' in which [*the* first memes arose. Once self-copying memes had arisen, their own, much faster, kind of evolution took off. We biologists have assimilated the idea of genetic evolution so deeply that we tend to forget that it is only one of many possible kinds of evolution.

Imitation, in the broad sense, is how memes *can* replicate. But just as not all genes that can replicate do so successfully, so some memes are more successful in the meme-pool than others. This is the analogue of natural selection. I have mentioned particular examples of qualities that make for high survival value among memes. But in general they must be the same as those discussed for the replicators of Chapter 2: longevity, fecundity, and copying-fidelity. The longevity of any one copy of a meme is probably relatively unimportant, as it is for any one copy of a gene. The copy of the tune 'Auld Lang Syne' that exists in my brain will last only for the rest of my life.* The copy of the same tune that is printed in my volume of *The Scottish Student's Song Book* is unlikely to last much longer. But I expect there will be copies of the same tune on paper and in peoples' brains for centuries to come. As in the case of genes, fecundity is much more important than longevity of particular copies. If the meme is a scientific idea, its spread will depend on how acceptable it is to the population of individual scientists; a rough measure of its survival value could be obtained by counting the number of times it is referred to in

successive years in scientific journals.* If it is a popular tune, its spread through the meme pool may be gauged by the number of people heard whistling it in the streets. If it is a style of women's shoe, the population memeticist may use sales statistics from shoe shops. Some memes, like some genes, achieve brilliant short-term success in spreading rapidly, but do not last long in the meme pool. Popular songs and stiletto heels are examples. Others, such as the Jewish religious laws, may continue to propagate themselves for thousands of years, usually because of the great potential permanence of written records.

This brings me to the third general quality of successful replicators: copying-fidelity. Here I must admit that I am on shaky ground. At first sight it looks as if memes are not high-fidelity replicators at all. Every time a scientist hears an idea and passes it on to somebody else, he is likely to change it somewhat. I have made no secret of my debt in this book to the ideas of R. L. Trivers. Yet I have not repeated them in his own words. I have twisted them round for my own purposes, changing the emphasis, blending them with ideas of my own and of other people. The memes are being passed on to you in altered form. This looks quite unlike the particulate, all-or-none quality of gene transmission. It looks as though meme transmission is subject to continuous mutation, and also to blending.

It is possible that this appearance of non-particulate nature is illusory, and that the analogy with genes does not break down. After all, if we look at the inheritance of many genetic characters such as human height or skin-colouring, it does not look like the work of indivisible and unblendable genes. If a black and a white person mate, their children do not come out either black or white: they are intermediate. This does not mean the genes concerned are not particulate. It is just that there are so many of them concerned with skin colour, each one having such a small effect, that they *seem to blend*. So far I have talked of memes as though it was obvious what a single unit-meme consisted of. But of course it is far from obvious. I have said a tune is one meme, but what about a symphony: how many memes is that? Is each movement one meme, each recognizable phrase of melody, each bar, each chord, or what?

I appeal to the same verbal trick as I used in Chapter 3. There I divided the 'gene complex' into large and small genetic units, and units within units. The 'gene' was defined, not in a rigid all-or-none way, but as a unit of convenience, a length of chromosome with just sufficient copying-fidelity to serve as a viable unit of natural selection. If a single phrase of Beethoven's ninth symphony is sufficiently distinctive and memorable to be abstracted from the context of the whole symphony, and used as the call-sign of a maddeningly intrusive European broadcasting station, then to that extent it deserves to be called one meme. It has, incidentally, materially diminished my capacity to enjoy the original symphony.

Similarly, when we say that all biologists nowadays believe in Darwin's theory, we do not mean that every biologist has, graven in his brain, an identical copy of the exact words of Charles Darwin himself. Each individual has his own way of interpreting Darwin's ideas. He probably learned them not from Darwin's own writings, but from more recent authors. Much of what Darwin said is, in detail, wrong. Darwin if he read this book would scarcely recognize his own original theory in it, though I hope he would like the way I put it. Yet, in spite of all this, there is something, some essence of Darwinism, which is present in the head of every individual who understands the theory. If this were not so, then almost any statement about two people agreeing with each other would be meaningless. An 'idea-meme' might be defined as an entity that is capable of being transmitted from one brain to another. The meme of Darwin's theory is therefore that essential basis of the idea which is held in common by all brains that understand the theory. The *differences* in the ways that people represent the theory are then, by definition, not part of the meme. If Darwin's theory can be subdivided into components, such that some people believe component *A* but not component *B*, while others believe *B* but *not A*, then *A* and *B* should be regarded as separate memes. If almost everybody who believes in *A* also believes in *B*—if the memes are closely 'linked' to use the genetic term—then it is convenient to lump them together as one meme.

Let us pursue the analogy between memes and genes further. Throughout this book, I have emphasized that we must not think of genes as conscious, purposeful agents. Blind natural selection, however, makes them behave rather as if they were purposeful, and it has been convenient, as a

shorthand, to refer to genes in the language of purpose. For example, when we say 'genes are trying to increase their numbers in future gene pools', what we really mean is 'those genes that behave in such a way as to increase their numbers in future gene pools tend to be the genes whose effects we see in the world'. Just as we have found it convenient to think of genes as active agents, working purposefully for their own survival, perhaps it might be convenient to think of memes in the same way. In neither case must we get mystical about it. In both cases the idea of purpose is only a metaphor, but we have already seen what a fruitful metaphor it is in the case of genes. We have even used words like 'selfish' and 'ruthless' of genes, knowing full well it is only a figure of speech. Can we, in exactly the same spirit, look for selfish or ruthless memes?

There is a problem here concerning the nature of competition. Where there is sexual reproduction, each gene is competing particularly with its own alleles—rivals for the same chromosomal slot. Memes seem to have nothing equivalent to chromosomes, and nothing equivalent to alleles. I suppose there is a trivial sense in which many ideas can be said to have 'opposites'. But in general memes resemble the early replicating molecules, floating chaotically free in the primeval soup, rather than modern genes in their neatly paired, chromosomal regiments. In what sense then are memes competing with each other? Should we expect them to be 'selfish' or 'ruthless', if they have no alleles? The answer is that we might, because there is a sense in which they must indulge in a kind of competition with each other.

Any user of a digital computer knows how precious computer time and memory storage space are. At many large computer centers they are literally costed in money; or each user may be allotted a ration of time, measured in seconds, and a ration of space, measured in 'words'. The computers in which memes live are human brains.* Time is possibly a more important limiting factor than storage space, and it is the subject of heavy competition. The human brain, and the body that it controls, cannot do more than one or a few things at once. If a meme is to dominate the attention of a human brain, it must do so at the expense of 'rival' memes. Other commodities for which memes compete are radio and television time, billboard space, newspaper column-inches, and library shelf-space.

In the case of genes, we saw in Chapter 3 that co-adapted gene complexes may arise in the gene pool. A large set of genes concerned with mimicry in butterflies became tightly linked together on the same chromosome, so tightly that they can be treated as one gene. In Chapter 5 we met the more sophisticated idea of the evolutionarily stable set of genes. Mutually suitable teeth, claws, guts, and sense organs evolved in carnivore gene pools, while a different stable set of characteristics emerged from herbivore gene pools. Does anything analogous occur in meme pools? Has the god meme, say, become associated with any other particular memes, and does this association assist the survival of each of the participating memes? Perhaps we could regard an organized church, with its architecture, rituals, laws, music, art, and written tradition, as a co-adapted stable set of mutually-assisting memes.

To take a particular example, an aspect of doctrine that has been very effective in enforcing religious observance is the threat of hell fire. Many children and even some adults believe that they will suffer ghastly torments after death if they do not obey the priestly rules. This is a peculiarly nasty technique of persuasion, causing great psychological anguish throughout the middle ages and even today. But it is highly effective. It might almost have been planned deliberately by a machiavellian priesthood trained in deep psychological indoctrination techniques. However, I doubt if the priests were that clever. Much more probably, unconscious memes have ensured their own survival by virtue of those same qualities of pseudo-ruthlessness that successful genes display. The idea of hell fire is, quite simply, *self perpetuating*, because of its own deep psychological impact. It has become linked with the god meme because the two reinforce each other, and assist each other's survival in the meme pool.

Another member of the religious meme complex is called faith. It means blind trust, in the absence of evidence, even in the teeth of evidence. The story of Doubting Thomas is told, not so that we shall admire Thomas, but so that we can admire the other apostles in comparison. Thomas demanded evidence. Nothing is more lethal for certain kinds of meme than a tendency to look for evidence. The other apostles,

whose faith was so strong that they did not need evidence, are held up to us as worthy of imitation. The meme for blind faith secures its own perpetuation by the simple unconscious expedient of discouraging rational inquiry.

Blind faith can justify anything.* If a man believes in a different god, or even if he uses a different ritual for worshipping the same god, blind faith can decree that he should die—on the cross, at the stake, skewered on a Crusader's sword, shot in a Beirut street, or blown up in a bar in Belfast. Memes for blind faith have their own ruthless ways of propagating themselves. This is true of patriotic and political as well as religious blind faith.

Memes and genes may often reinforce each other, but they sometimes come into opposition. For example, the habit of celibacy is presumably not inherited genetically. A gene for celibacy is doomed to failure in the gene pool, except under very special circumstances such as we find in the social insects. But still, a *meme* for celibacy can be successful in the meme pool. For example, suppose the success of a meme depends critically on how much time people spend in actively transmitting it to other people. Any time spent in doing other things than attempting to transmit the meme may be regarded as time wasted from the meme's point of view. The meme for celibacy is transmitted by priests to young boys who have not yet decided what they want to do with their lives. The medium of transmission is human influence of various kinds, the spoken and written word, personal example and so on. Suppose, for the sake of argument, it happened to be the case that marriage weakened the power of a priest to influence his flock, say because it occupied a large proportion of his time and attention. This has, indeed, been advanced as an official reason for the enforcement of celibacy among priests. If this were the case, it would follow that the meme for celibacy could have greater survival value than the meme for marriage. Of course, exactly the opposite would be true for a *gene* for celibacy. If a priest is a survival machine for memes, celibacy is a useful attribute to build into him. Celibacy is just a minor partner in a large complex of mutually-assisting religious memes.

I conjecture that co-adapted meme-complexes evolve in the same kind of way as co-adapted gene-complexes. Selection favours memes that exploit their cultural environment to their own advantage. This cultural environment consists of other memes which are also being selected. The meme pool therefore comes to have the attributes of an evolutionarily stable set, which new memes find it hard to invade.

I have been a bit negative about memes, but they have their cheerful side as well. When we die there are two things we can leave behind us: genes and memes. We were built as gene machines, created to pass on our genes. But that aspect of us will be forgotten in three generations. Your child, even your grandchild, may bear a resemblance to you, perhaps in facial features, in a talent for music, in the colour of her hair. But as each generation passes, the contribution of your genes is halved. It does not take long to reach negligible proportions. Our genes may be immortal but the *collection* of genes that is any one of us is bound to crumble away. Elizabeth II is a direct descendant of William the Conqueror. Yet it is quite probable that she bears not a single one of the old king's genes. We should not seek immortality in reproduction.

But if you contribute to the world's culture, if you have a good idea, compose a tune, invent a sparking plug, write a poem, it may live on, intact, long after your genes have dissolved in the common pool. Socrates may or may not have a gene or two alive in the world today, as G. C. Williams has remarked, but who cares? The meme-complexes of Socrates, Leonardo, Copernicus and Marconi are still going strong.

However speculative my development of the theory of memes may be, there is one serious point which I would like to emphasize once again. This is that when we look at the evolution of cultural traits and at their survival value, we must be clear *whose* survival we are talking about. Biologists, as we have seen, are accustomed to looking for advantages at the gene level (or the individual, the group, or the species level according to taste). What we have not previously considered is that a cultural trait may have evolved in the way that it has, simply because it is *advantageous to itself*.

We do not have to look for conventional biological survival values of traits like religion, music, and ritual dancing, though these may also be present. Once the genes have provided their survival machines with brains that are capable of rapid imitation, the memes will automatically take over. We do not even have to posit a genetic advantage in imitation, though that would certainly help. All that is necessary is that the brain should be *capable* of imitation: memes will then evolve that exploit the capability to the full.

I now close the topic of the new replicators, and end the chapter on a note of qualified hope. One unique feature of man, which may or may not have evolved memically, is his capacity for conscious foresight. Selfish genes (and, if you allow the speculation of this chapter, memes too) have no foresight. They are unconscious, blind, replicators. The fact that they replicate, together with certain further conditions means, willy nilly, that they will tend towards the evolution of qualities which, in the special sense of this book, can be called selfish. A simple replicator, whether gene or meme, cannot be expected to forgo short-term selfish advantage even if it would really pay it, in the long term, to do so. We saw this in the chapter on aggression. Even though a 'conspiracy of doves' would be better for *every single individual* than the evolutionarily stable strategy, natural selection is bound to favour the ESS.

It is possible that yet another unique quality of man is a capacity for genuine, disinterested, true altruism. I hope so, but I am not going to argue the case one way or the other, nor to speculate over its possible memic evolution. The point I am making now is that, even if we look on the dark side and assume that individual man is fundamentally selfish, our conscious foresight—our capacity to simulate the future in imagination—could save us from the worst selfish excesses of the blind replicators. We have at least the mental equipment to foster our long-term selfish interests rather than merely our short-term selfish interests. We can see the long-term benefits of participating in a 'conspiracy of doves', and we can sit down together to discuss ways of making the conspiracy work. We have the power to defy the selfish genes of our birth and, if necessary, the selfish memes of our indoctrination. We can even discuss ways of deliberately cultivating and nurturing pure, disinterested altruism— something that has no place in nature, something that has never existed before in the whole history of the world. We are built as gene machines and cultured as meme machines, but we have the power to turn against our creators. We, alone on earth, can rebel against the tyranny of the selfish replicators.*

Mememes: the new replicators

p. 192 *I would put my money on one fundamental principle... all life evolves by the differential survival of replicating entities.*

My wager that all life, everywhere in the universe, would turn out to have evolved by Darwinian means has now been spelled out and justified more fully in my paper 'Universal Darwinism' and in the last chapter of *The Blind Watchmaker*. I show that all the alternatives to Darwinism that have ever been suggested are in principle incapable of doing the job of explaining the organized complexity of life. The argument is a general one, not based upon particular facts about life as we know it. As such it has been criticized by scientists pedestrian enough to think that slaving over a hot test tube (or cold muddy boot) is the only method of discovery in science. One critic complained that my argument was 'philosophical', as though that was sufficient condemnation. Philosophical or not, the fact is that neither he nor anybody else has found any flaw in what I said. And 'in principle' arguments such as mine, far from being irrelevant to the real world, can be *more* powerful than arguments based on particular factual research. My reasoning, if it is correct, tells us something important about life everywhere in the universe. Laboratory and field research can tell us only about life as we have sampled it here.

p. 192 *Meme*

The word meme seems to be turning out to be a good meme. It is now quite widely used and in 1988 it joined the official list of words being considered for future editions of Oxford English Dictionaries. This makes me the more anxious to repeat that my designs on human culture were modest almost to vanishing point. My true ambitions—and they are admittedly large—lead in another direction entirely. I want to claim almost limitless power for slightly inaccurate self-replicating entities, once they arise anywhere in the universe. This is because they tend to become the basis for Darwinian selection which, given enough generations, cumulatively builds systems of great complexity. I believe that, given the right conditions, replicators automatically band together to create systems, or machines, that carry them around and work to favour their continued replication. The first ten chapters of *The Selfish Gene* had concentrated exclusively on one kind of replicator, the gene. In discussing mememes in the final chapter I was trying to make the case for replicators in general, and to show that genes were not the only members of that important class. Whether the milieu of human culture really does have what it takes to get a form of Darwinism going, I am not sure. But in any case that question is subsidiary to my concern. Chapter 11 will have succeeded if the reader closes the book with the feeling that DNA molecules are not the only entities that might form the basis for Darwinian evolution. My purpose was to cut the gene down to size, rather than to sculpt a grand theory of human culture.

p. 192 ... *mememes should be regarded as living structures, not just metaphorically but technically.*

DNA is a self-replicating piece of hardware. Each piece has a particular structure, which is different from rival pieces of DNA. If mememes in brains are analogous to genes they must be self-replicating brain structures, actual patterns of neuronal wiring-up that reconstitute themselves in one brain after another. I had always felt uneasy spelling this out aloud, because we know far less about brains than about genes, and are therefore necessarily vague about what such a brain structure might actually be. So I was relieved to receive recently a very interesting paper by Juan Delius of the University of Konstanz in Germany. Unlike me, Delius doesn't have to feel apologetic, because he is a distinguished brain scientist whereas I am not a brain scientist at all. I am delighted, therefore, that he is bold enough to ram home the point by actually publishing a detailed picture of what the neuronal hardware of a meme might look like. Among the other interesting things he does is to explore, far more searchingly than I had done, the analogy of mememes with parasites; to be more precise, with the spectrum of which malignant parasites are one extreme, benign 'symbionts' the other extreme. I am particularly keen on this approach because of

my own interest in 'extended phenotypic' effects of parasite genes on host behaviour (see Chapter 13 of this book and especially chapter 12 of *The Extended Phenotype*). Delius, by the way, emphasizes the clear separation between memes and their ('phenotypic') effects. And he reiterates the importance of co-adapted meme-complexes, in which memes are selected for their mutual compatibility.

p. 194 '*Auld Long Syne*'

'Auld Lang Syne' was, unwittingly, a revealingly fortunate example for me to have chosen. This is because, almost universally, it is rendered with an error, a mutation. The refrain is, essentially always nowadays, sung as 'For the sake of auld lang syne', whereas Burns actually wrote 'For auld lang syne'. A memically minded Darwinian immediately wonders what has been the 'survival value' of the interpolated phrase, 'the sake of'. Remember that we are not looking for ways in which *people* might have survived better through singing the song in altered form. We are looking for ways in which the alteration *itself* might have been good at surviving in the meme pool. Everybody learns the song in childhood, not through reading Burns but through hearing it sung on New Year's Eve. Once upon a time, presumably, everybody sang the correct words. 'For the sake of' must have arisen as a rare mutation. Our question is, why has the initially rare mutation spread so insidiously that it has become the norm in the meme pool?

I don't think the answer is far to seek. The sibilant 's' is notoriously obtrusive. Church choirs are drilled to pronounce 's' sounds as lightly as possible, otherwise the whole church echoes with hissing. A murmuring priest at the altar of a great cathedral can sometimes be heard, from the back of the nave, only as a sporadic sussuration of 's's. The other consonant in 'sake', 'k', is almost as penetrating. Imagine that nineteen people are correctly singing 'For auld lang syne' and one person, somewhere in the room, slips in the erroneous 'For the sake of auld lang syne'. A child, hearing the song for the first time, is eager to join in but uncertain of the words. Although almost everybody is singing 'For auld lang syne', the hiss of an 's' and the cut of a 'k' force their way into the child's ears, and when the refrain comes round again he too sings 'For the sake of auld lang syne'. The mutant meme has taken over another vehicle. If there are any other children there, or adults unconfident of the words, they will be more likely to switch to the mutant form next time the refrain comes round. It is not that they 'prefer' the mutant form. They genuinely don't know the words and are honestly eager to learn them. Even if those who know better indignantly bellow 'For auld lang syne' at the top of their voice (as I do!), the correct words happen to have no conspicuous consonants, and the mutant form, even if quietly and diffidently sung, is far easier to hear.

A similar case is 'Rule Britannia'. The correct second line of the chorus is 'Britannia, rule the waves.' It is frequently, though not quite universally, sung as 'Britannia rules the waves.' Here the insistently hissing 's' of the meme is aided by an additional factor. The intended meaning of the poet (James Thomson) was presumably imperative (Britannia, go out and rule the waves!) or possibly subjunctive (let Britannia rule the waves). But it is superficially easier to misunderstand the sentence as indicative (Britannia, as a matter of fact, does rule the waves). This mutant meme, then, has two separate survival values over the original form that it replaced: it sounds more conspicuous and it is easier to understand.

The final test of a hypothesis should be experimental. It should be possible to inject the hissing meme, deliberately, into the meme pool at a very low frequency, and then watch it spread because of its own survival value. What if just a few of us were to start singing 'God saves our gracious Queen'?

p. 194 If the meme is a scientific idea, its spread will depend on how acceptable it is to the population of individual scientists; a rough measure of its survival value could be obtained by counting the number of times it is referred to in successive years in scientific journals.

I'd hate it if this were taken to mean that 'catchiness' was the only criterion for acceptance of a scientific idea. After all, some scientific ideas are actually right, others wrong! Their rightness and wrongness can be tested; their logic can be dissected. They are really not like pop-tunes, religions, or punk hairdos. Nevertheless there is a sociology as well as a logic to science. Some bad scientific ideas can spread widely, at least for a while. And some good ideas lie dormant for years before finally catching on and

colonizing scientific imaginations.

We can find a prime example of this dormancy followed by rampant propagation in one of the main ideas in this book, Hamilton's theory of kin selection. I thought it would be a fitting case for trying out the idea of measuring meme spread by counting journal references. In the first edition I noted (p. 90) that 'His two papers of 1964 are among the most important contributions to social ethology ever written, and I have never been able to understand why they have been so neglected by ethologists (his name does not even appear in the index of two major text-books of ethology, both published in 1970). Fortunately there are recent signs of a revival of interest in his ideas.' I wrote that in 1976. Let us trace the course of that memic revival over the subsequent decade.

Science Citation Index is a rather strange publication in which one may look up any published paper and see tabulated, for a given year, the number of subsequent publications that have quoted it. It is intended as an aid to tracking down the literature on a given topic. University appointments committees have picked up the habit of using it as a rough and ready (too rough and too ready) way of comparing the scientific achievements of applicants for jobs. By counting the citations of Hamilton's papers, in each year since 1964, we can approximately track the progress of his ideas into the consciousness of biologists (Figure i). The initial dormancy is very evident. Then it looks as though there is a dramatic upturn in interest in kin selection during the 1970s. If there is any point where the upward trend begins, it seems to be between 1973 and 1974. The upturn then gathers pace up to a peak in 1981, after which the annual rate of citation fluctuates irregularly about a plateau.

A memic myth has grown up that the upsurge of interest in kin selection was all triggered by books published in 1975 and 1976. The graph, with its upturn in 1974, seems to give the lie to this idea. On the contrary, the evidence could be used to support a very different hypothesis, namely that we are dealing with one of those ideas that was 'in the air', 'whose time had come'. Those mid-seventies books would, on this view, be symptoms of the bandwagon effect rather than prime causes of it.

[I'll bring in copies of these diagrams – LHN]

Perhaps, indeed, we are dealing with a longer-term, slow-starting, exponentially accelerating bandwagon that began much earlier. One way of testing this simple, exponential hypothesis is to plot the citations cumulatively on a *logarithmic* scale. Any growth process, where rate of growth is proportional to size already attained, is called exponential growth. A typical exponential process is an epidemic: each person breathes the virus on several other people, each of whom in turn breathes on the same number again, so the number of victims grows at an ever increasing rate. It is diagnostic of an exponential curve that it becomes a straight line when plotted on a logarithmic scale. It is not necessary, but it is convenient and conventional, to plot such logarithmic graphs cumulatively. If the spread of Hamilton's meme was really like a gathering epidemic, the points on a cumulative logarithmic graph should fall on a single straight line. Do they?

The particular line drawn in Figure 2 is the straight line that, statistically speaking, best fits all the points. The apparent sharp rise between 1966 and 1967 should probably be ignored as an unreliable small-numbers effect of the kind that logarithmic plotting would tend to exaggerate. Thereafter, the graph is not a bad approximation to a single straight line, although minor overlying patterns can also be discerned. If my exponential interpretation is accepted, what we are dealing with is a single slow-burning explosion of interest, running right through from 1967 to the late 1980s. Individual books and papers should be seen both as symptoms and as causes of this long-term trend.

Do not think by the way, that this pattern of increase is somehow trivial, in the sense of being inevitable. Any cumulative curve would, of course, rise even if the rate of citations per year were constant. But on the logarithmic scale it would rise at a steadily slower rate: it would tail off. The thick line at the top of Figure 3 shows the *theoretical* curve that we would get if every year had a constant citation rate (equal to the actual average rate of Hamilton citations, of about 37 per year). This dying away *curve* can be compared

directly with the observed *straight* line in Figure 2. which indicates an exponential rate of increase. We really do have a case of increase upon increase, not a steady rate of citation.

FIGURE 3. Log cumulative citations of three works not by Hamilton-compared with 'theoretical' curve for Hamilton (details explained in text)

Secondly, one might be tempted to think that there is something, if not inevitable, at least trivially expected about an exponential increase. Isn't the whole rate of publication of scientific papers, and therefore opportunities to cite other papers, itself increasing exponentially? Perhaps the size of the scientific community is increasing exponentially. To show that there is something special about the Hamilton meme, the easiest way is to plot the same kind of graph for some other papers. Figure 3 also shows the log cumulative citation frequencies of three other works (which incidentally were also highly influential on the first edition of this book). These are Williams's 1966 book, *Adaptation and Natural Selection*; Trivers's 1971 paper on reciprocal altruism; and Maynard Smith and Price's 1973 paper introducing the ESS idea. All three of them show curves that clearly are not exponential over the whole time-span. For these works too, however, the annual citation rates are far from uniform, and over part of their range they may even be exponential. The Williams curve, for instance, is approximately a straight line on the log scale from about 1970 onwards, suggesting that it, too entered an explosive phase of influence.

I have been downplaying the influence of particular books in spreading the Hamilton meme. Nevertheless, there is one apparently suggestive postscript to this little piece of memic analysis. As in the case of 'Auld Lang Syne' and 'Rule Britannia', we have an illuminating mutant error. The correct title of Hamilton's 1964 pair of papers was 'The genetical evolution of social behaviour'. In the mid to late 1970s, a rash of publications, *Sociobiology* and *The Selfish Gene* among them, mistakenly cited it as 'The genetical theory of social behaviour'. Jon Seger and Paul Harvey looked for the earliest occurrence of this mutant meme, thinking that it would be a neat marker, almost like a radioactive label, for tracing scientific influence. They traced it back to E.O. Wilson's influential book, *Sociobiology*, published in 1975, and even found some indirect evidence for this suggested pedigree.

Much as I admire Wilson's *tour de force*—I wish people would read it more and read about it less—my hackles have always risen at the entirely false suggestion that his book influenced mine. Yet, since my book also contained the mutant citation—the 'radioactive label'—it began to look alarmingly as though at least one meme had travelled from Wilson to me! This would not have been particularly surprising, since *Sociobiology* arrived in Britain just as I was completing *The Selfish Gene*, the very time when I would have been working on my bibliography. Wilson's massive bibliography would have seemed a godsend, saving hours in the library. My chagrin turned to glee, therefore, when I chanced upon an old stencilled bibliography that I had handed to the students at an Oxford lecture in 1970. Large as life, there was 'The genetical *theory* of social behaviour', a whole five years earlier than Wilson's publication. Wilson couldn't possibly have seen my 1970 bibliography. There was no doubt about it: Wilson and I had independently introduced the same mutant meme!

How could such a coincidence have happened? Once again, as in the case of 'Auld Lang Syne', a plausible explanation is not far to seek. R. A. Fisher's most famous book is called *The Genetical Theory of Natural Selection*. Such a household name has this title become in the world of evolutionary biologists, it is hard for us to hear its first two words without automatically adding the third. I suspect that both Wilson and I must have done just that. This is a happy conclusion for all concerned, since nobody minds admitting to being influenced by Fisher!

p. 197 *The computers in which memes live are human brains.*

It was obviously predictable that manufactured electronic computers, too, would eventually play host to self-replicating patterns of information— memes. Computers are increasingly tied together in intricate networks of shared information. Many of them are literally wired up together in electronic mail exchange.

Others share information when their owners pass floppy discs around. It is a perfect milieu for self-replicating programs to flourish and spread. When I wrote the first edition of this book I was naive enough to suppose that an undesirable computer meme would have to arise by a spontaneous error in the copying of a legitimate program, and I thought this an unlikely event. Alas, that was a time of innocence. Epidemics of 'viruses' and 'worms', deliberately released by malicious programmers, are now familiar hazards to computer-users all over the world. My own hard disc has to my knowledge been infected in two different virus epidemics during the past year, and that is a fairly typical experience among heavy computer-users. I shall not mention the names of particular viruses for fear of giving any nasty little satisfaction to their nasty little perpetrators. I say 'nasty', because their behaviour seems to me morally indistinguishable from that of a technician in a microbiology laboratory, who deliberately infects the drinking water and seeds epidemics in order to snigger at people getting ill. I say 'little', because these people are mentally little. There is nothing clever about designing a computer virus. Any half-way competent programmer could do it, and half-way competent programmers are two-a-penny in the modern world. I'm one myself. I shan't even bother to explain how computer viruses work. It's only too obvious.

What is less easy is to know how to combat them. Unfortunately some very expert programmers have had to waste their valuable time writing virus-detector programs, immunization programs and so on (the analogy with medical vaccination, by the way, is astonishingly close, even down to the injection of a 'weakened strain' of the virus). The danger is that an arms race will develop, with each advance in virus-prevention being matched by counter-advances in new virus programs. So far, most anti-virus programs are written by altruists and supplied free of charge as a service. But I foresee the growth of a whole new profession—splitting into lucrative specialisms just like any other profession—of 'software doctors', on call with black bags full of diagnostic and curative floppy discs. I use the name 'doctors', but real doctors are solving natural problems that are not deliberately engineered by human malice. My software doctors, on the other hand, will be, like lawyers, solving man-made problems that should never have existed in the first place. In so far as virus-makers have any discernible motive, they presumably feel vaguely anarchistic. I appeal to them: do you really want to pave the way for a new fat-cat profession? If not, stop playing at silly memes, and put your modest programming talents to better use.

p. 198 *Blind faith can justify anything.*

I have had the predictable spate of letters from faith's victims, protesting about my criticisms of it. Faith is such a successful brainwasher in its own favour, especially a brainwasher of children, that it is hard to break its hold. But what, after all, is faith? It is a state of mind that leads people to believe something—it doesn't matter what—in the total absence of supporting evidence. If there were good supporting evidence then faith would be superfluous, for the evidence would compel us to believe it anyway. It is this that makes the often-parroted claim that 'evolution itself is a matter of faith' so silly. People believe in evolution not because they arbitrarily want to believe it but because of overwhelming, publicly available evidence.

I said 'it doesn't matter what' the faithful believe, which suggests that people have faith in entirely daft, arbitrary things, like the electric monk in Douglas Adams's delightful *Dirk Gently's Holistic Detective Agency*. He was purpose-built to do your believing for you, and very successful at it. On the day that we meet him he unshakeably believes, against all the evidence, that everything in the world is pink. I don't want to argue that the things in which a particular individual has faith are necessarily daft. They may or may not be. The point is that there is no way of deciding whether they are, and no way of preferring one article of faith over another, because evidence is explicitly eschewed. Indeed the fact that true faith doesn't need evidence is held up as its greatest virtue; this was the point of my quoting the story of Doubting Thomas, the only really admirable member of the twelve apostles.

Faith cannot move mountains (though generations of children are solemnly told the contrary and believe it). But it is capable of driving people to such dangerous folly that faith seems to me to qualify as a kind of mental illness. It leads people to believe in whatever it is so strongly that in extreme cases they are prepared to kill and to die for it without the need for further justification. Keith Henson has coined the

name 'memeoids' for 'victims that have been taken over by a meme to the extent that their own survival becomes inconsequential... You see lots of these people on the evening news from such places as Belfast or Beirut.' Faith is powerful enough to immunize people against all appeals to pity, to forgiveness, to decent human feelings. It even immunizes them against fear, if they honestly believe that a martyr's death will send them straight to heaven. What a weapon! Religious faith deserves a chapter to itself in the annals of war technology, on an even footing with the longbow, the warhorse, the tank, and the hydrogen bomb.

p. 201 *We, alone on earth, can rebel against the tyranny of the selfish replicators.*

The optimistic tone of my conclusion has provoked skepticism among critics who feel that it is inconsistent with the rest of the book. In some cases the criticism comes from doctrinaire sociobiologists jealously protective of the importance of generic influence. In other cases the criticism comes from a paradoxically opposite quarter, high priests of the left jealously protective of a favourite demonological icon! Rose, Kamin, and Lewontin in *Not in Our Genes* have a private bogey called 'reductionism'; and all the best reductionists are also supposed to be 'determinists', preferably 'genetic determinists'.

Brains, for reductionists, are determinate biological objects whose properties produce the behaviors we observe and the states of thought or intention we infer from that behavior ... Such a position is, or ought to be, completely in accord with the principles of sociobiology offered by Wilson and Dawkins. However, to adopt it would involve them in the dilemma of first arguing the innateness of much human behavior that, being liberal men, they clearly find unattractive (spite, indoctrination, etc.) and then to become entangled in liberal ethical concerns about responsibility for criminal acts, if these, like all other acts, are biologically determined. To avoid this problem, Wilson and Dawkins invoke a free will that enables us to go against the dictates of our genes if we so wish... This is essentially a return to unabashed Cartesianism, a dualistic *deus ex machina*.

I *think* that Rose and his colleagues are accusing us of eating our cake and having it. Either we must be 'genetic determinists' or we believe in 'free will'; we cannot have it both ways. But—and here I presume to speak for Professor Wilson as well as for myself—it is only in the eyes of Rose and his colleagues that we are 'genetic determinists'. What they don't understand (apparently, though it is hard to credit) is that it is perfectly possible to hold that genes exert a statistical influence on human behaviour while at the same time believing that this influence can be modified, overridden or reversed by other influences. Genes must exert a statistical influence on any behaviour pattern that evolves by natural selection. Presumably Rose and his colleagues agree that human sexual desire has evolved by natural selection, in the same sense as anything ever evolves by natural selection. They therefore must agree that there have been genes influencing sexual desire—in the same sense as genes ever influence anything. Yet they presumably have no trouble with curbing their sexual desires when it is socially necessary to do so. What is dualist about that? Obviously nothing. And no more is it dualist for me to advocate rebelling 'against the tyranny of the selfish replicators'. We, that is, our brains, are separate and independent enough from our genes to rebel against them. As already noted, we do so in a small way every time we use contraception. There is no reason why we should not rebel in a large way, too.