Foreword

by Richard Dawkins

As an undergraduate I was chatting to a friend in the Balliol College lunch queue. He regarded me with increasingly quizzical amusement, then asked: 'Have you just been with Peter Brunet?' I had indeed, though I couldn't guess how he knew. Peter Brunet was our much loved tutor, and I had come hotfoot from a tutorial hour with him. 'I thought so', my friend laughed. 'You are talking just like him; your voice sounds exactly like his.' I had, if only briefly, 'inherited' intonations and manners of speech from an admired, and now greatly missed, teacher. Years later, when I became a tutor myself, I taught a young woman who affected an unusual habit. When asked a question which required deep thought, she would screw her eyes tight shut, jerk her head down to her chest and then freeze for up to half a minute before looking up, opening her eyes, and answering the question with fluency and intelligence. I was amused by this, and did an imitation of it to divert my colleagues after dinner. Among them was a distinguished Oxford philosopher. As soon as he saw my imitation, he immediately said: 'That's Wittgenstein! Is her surname ____ by any chance?' Taken aback, I said that it was. 'I thought so', said my colleague. 'Both her parents are professional philosophers and devoted followers of Wittgenstein.' The gesture had passed from the great philosopher, via one or both of her parents to my pupil. I suppose that, although my further imitation was done in jest, I must count myself a fourth-generation transmitter of the gesture. And who knows where Wittgenstein got it?

The fact that we unconsciously imitate others, especially our parents, those in quasi-parental roles, or those we admire, is familiar enough. But is it really credible that imitation could become the basis of a major theory of the evolution of the human mind and the explosive inflation of the human brain, even of what it means to be a conscious self? Could imitation have been the key to what set our ancestors apart from all other animals? I would never have thought so, but Susan Blackmore in this book makes a tantalisingly strong case.

Imitation is how a child learns its particular language rather than some

other language. It is why people speak more like their own parents than like other people's parents. It is why regional accents, and on a longer timescale separate languages, exist. It is why religions persist along family lines rather than being chosen afresh in every generation. There is at least a superficial analogy to the longitudinal transmission of genes down generations, and to the horizontal transmission of genes in viruses. Without prejudging the issue of whether the analogy is a fruitful one, if we want even to talk about it we had better have a name for the entity that might play the role of gene in the transmission of words, ideas, faiths, mannerisms and fashions. Since 1976, when the word was coined, increasing numbers of people have adopted the name 'meme' for the postulated gene analogue.

The compilers of the *Oxford English Dictionary* operate a sensible criterion for deciding whether a new word shall be canonised by inclusion. The aspirant word must be commonly used without needing to be defined and without its coinage being attributed whenever it is used. To ask the metamemetic question, how widespread is 'meme'? A far from ideal, but nevertheless easy and convenient method of sampling the meme pool, is provided by the World Wide Web and the ease with which it may be searched. I did a quick search of the Web on the day of writing this, which happened to be 29 August 1998. 'Meme' is mentioned about half a million times, but that is a ridiculously high figure, obviously confounded by various acronyms and the French *même*. The adjectival form 'memetic', however, is genuinely exclusive, and it clocked up 5042 mentions. To put this number into perspective, I compared a few other recently coined words or fashionable expressions. Spin doctor (or spin-doctor) gets 1412 mentions, dumbing down 3905, docudrama (or docu-drama) 2848, sociobiology 6679, catastrophe theory 1472, edge of chaos 2673, wannabee 2650, zippergate 1752, studmuffin 776, post-structural (or poststructural) 577, extended phenotype 515, exaptation 307. Of the 5042 mentions of memetic, more than 90 per cent make no mention of the origin of the word, which suggests that it does indeed meet the *OED*'s criterion. And, as Susan Blackmore tells us, the *Oxford English Dictionary* now does contain the following definition:

> meme An element of a culture that may be considered to be passed on by non-genetic means, esp. imitation.

Further searching of the Internet reveals a newsgroup talking shop, 'alt.memetics', which has received about 12 000 postings during the past year. There are on-line articles on, among many other things, 'The New Meme', 'Meme, Counter-meme', 'Memetics: a Systems Metabiology', 'Memes, and Grinning Idiot Press', 'Memes, Metamemes and Politics', 'Cryonics, religions and memes', 'Selfish Memes and the evolution of cooperation', and 'Running down the Meme'. There are separate Web pages on 'Memetics', 'Memes', 'The C Memetic Nexus', 'Meme theorists on the Web', 'Meme of the week', 'Meme Central', 'Arkuat's Meme Workshop', 'Some pointers and a short introduction to memetics', 'Memetics Index' and 'Meme Gardening Page'. There is even a new religion (tongue-in-cheek, I *think*), called the 'Church of Virus', complete with its own list of Sins and Virtues, and its own patron saint (Saint Charles Darwin, canonised as 'perhaps the most influential memetic engineer of the modern era') and I was alarmed to discover a passing reference to 'St Dawkin'.

Susan Blackmore's book is preceded by two others entirely devoted to the subject of memes and both good in their different ways: Richard Brodie's *Virus of the Mind: The New Science of the Meme*, and Aaron Lynch's *Thought Contagion: How Belief Spreads through Society*. Most significant of all, the distinguished philosopher Daniel Dennett has adopted the idea of the meme, building it in as a cornerstone of his theory of mind, as developed in his two great books *Consciousness Explained*, and *Darwin's Dangerous Idea*.

Memes travel longitudinally down generations, but they travel horizontally too, like viruses in an epidemic. Indeed, it is largely horizontal epidemiology that we are studying when we measure the spread of words like 'memetic', 'docudrama' or 'studmuffin' over the Internet. Crazes among schoolchildren provide particularly tidy examples. When I was about nine, my father taught me to fold a square of paper to make an origami Chinese junk. It was a remarkable feat of artificial embryology, passing through a distinctive series of intermediate stages: catamaran with two hulls, cupboard with doors, picture in a frame, and finally the junk itself, fully seaworthy or at least bathworthy, complete with deep hold, and two flat decks each surmounted by a large, square-rigged sail. The point of the story is that I went back to school and infected my friends with the skill, and it then spread around the school with the speed of the measles and pretty much the same epidemiological time-course. I do not know whether the epidemic subsequently jumped to other schools (a boarding school is a somewhat isolated backwater of the meme pool). But I do know that my father himself originally picked up the Chinese Junk meme during an almost identical epidemic at the same school 25 years earlier. The earlier virus was launched by the school matron. Long after the old matron's departure, I had reintroduced her meme to a new cohort of small boys.

Before leaving the Chinese Junk, let me use it to make one more point. A favourite objection to the meme/gene analogy is that memes, if they exist at all, are transmitted with too low fidelity to perform a gene-like role in any realistically Darwinian selection process. The difference between high-fidelity genes and low-fidelity memes is assumed to follow from the fact that genes, but not memes, are digital. I am sure that the details of Wittgenstein's mannerism were far from faithfully reproduced when I imitated my pupil's imitation of her parents' imitation of Wittgenstein. The form and timing of the tic undoubtedly mutated over the generations, as in the childhood game of Chinese Whispers (Americans call it Telephone).

Suppose we assemble a line of children. A picture of, say, a Chinese junk is shown to the first child, who is asked to draw it. The drawing, but not the original picture, is then shown to the second child, who is asked to make her own drawing of it. The second child's drawing is shown to the third child, who draws it again, and so the series proceeds until the twentieth child, whose drawing is revealed to everyone and compared with the first. Without even doing the experiment, we know what the result will be. The twentieth drawing will be so unlike the first as to be unrecognisable. Presumably, if we lay the drawings out in order, we shall note some resemblance between each one and its immediate predecessor and successor, but the mutation rate will be so high as to destroy all semblance after a few generations. A trend will be visible as we walk from one end of the series of drawings to the other, and the direction of the trend will be degeneration. Evolutionary geneticists have long understood that natural selection cannot work unless the mutation rate is low. Indeed, the initial problem of overcoming the fidelity barrier has been described as the Catch-22 of the Origin of Life. Darwinism depends on high-fidelity gene replication. How then can the meme, with its apparently dismal lack of fidelity, serve as quasi-gene in any quasi-Darwinian process?

It is not always as dismal as you think and, as Susan Blackmore insists, high fidelity is not necessarily synonymous with digital. Suppose we set up our Chinese Whispers Chinese Junk game again, but this time with a crucial difference. Instead of asking the first child to copy a drawing of a junk, we teach her, by demonstration, to make an origami model of a junk. When she has mastered the skill and made her own junk, the first child is asked to turn round to the second child and teach him how to make one. So the skill passes down the line to the twentieth child. What will be the result of this experiment? What will the twentieth child produce, and what shall we observe if we lay the twenty efforts out in order along the ground? I have not done it, but I will make the following

confident prediction, assuming that we run the experiment many times on different groups of twenty children. In several of the experiments, a child somewhere along the line will forget some crucial step in the skill taught him by the previous child, and the line of phenotypes will suffer an abrupt macromutation which will presumably then be copied to the end of the line, or until another discrete mistake is made. The end result of such mutated lines will not bear any resemblance to a Chinese junk at all. But in a good number of experiments the skill will correctly pass all along the line, and the twentieth junk will be no worse and no better, on average, than the first junk. If we then lay the twenty junks out in order, some will be more perfect than others, but imperfections will not be copied on down the line. If the fifth child is ham-fisted and makes a clumsily asymmetrical or floppy junk, his quantitative errors will be corrected if the sixth child happens to be more dextrous. The twenty junks will not exhibit a progressive deterioration in the way that the twenty drawings of our first experiment undoubtedly would.

Why? What is the crucial difference between the two kinds of experiment? It is this: inheritance in the drawing experiment is Lamarckian (Blackmore calls it 'copying-the-product'). In the origami experiment it is Weismannian (Blackmore's 'copying-the-instructions'). In the drawing experiment, the phenotype in every generation is also the genotype – it is what is passed on to the next generation. In the origami experiment, what passes to the next generation is not the paper phenotype but a set of instructions for making it. Imperfections in the execution of the instructions result in imperfect junks (phenotypes) but they are not passed on to future generations: they are non-memetic. Here are the first five instructions in the Weismannian meme line of instructions for making a Chinese junk:

1. Take a square sheet of paper and fold all four corners exactly into the middle.
2. Take the reduced square so formed, and fold one side into the middle.
3. Fold the opposite side into the middle, symmetrically.
4. In the same way, take the rectangle so formed, and fold its two ends into the middle.
5. Take the small square so formed, and fold it backwards, exactly along the straight line where your last two folds met.

. . . and so on, through 20 or 30 instructions of this kind. These instructions, though I would not wish to call them digital, are potentially

of very high fidelity, just as if they were digital. This is because they all make reference to idealised tasks like 'fold the four corners exactly into the middle'. If the paper is not exactly square, or if a child folds ineptly so that, say, the first corner overshoots the middle and the fourth corner undershoots it, the junk that results will be inelegant. But the next child in the line will not copy the error, for she will assume that her instructor *intended* to fold all four corners into the exact centre of a perfect square. The instructions are self-normalising. The code is error-correcting. Plato would enjoy it: what passes down the line is an ideal essence of junk, of which each actual junk is an imperfect approximation.

The instructions are more effectively passed on if verbally reinforced, but they can be transmitted by demonstration alone. A Japanese child could teach an English one, though neither has a word of the other's language. In the same way, a Japanese master carpenter could convey his skills to an equally monoglot English apprentice. The apprentice would not copy obvious mistakes. If the master hit his thumb with a hammer, the apprentice would correctly guess, even without understanding the Japanese expletive '** **** **!', that he meant to hit the nail. He would not make a Lamarckian copy of the precise details of every hammer blow, but copy instead the inferred Weismannian instruction: drive the nail in with as many blows of your hammer as it takes your arm to achieve the same idealised end result as the master has achieved with his – a nail head flush with the wood.

I believe that these considerations greatly reduce, and probably remove altogether, the objection that memes are copied with insufficient high fidelity to be compared with genes. For me, the quasi-genetic inheritance of language, and of religious and traditional customs, teaches the same lesson. Another objection, discussed, like the first, in Susan Blackmore's illuminating chapter on 'Three problems with memes' is that we do not know what memes are made of or where they reside. Memes have not yet found their Watson and Crick; they even lack their Mendel. Where genes are to be found in precise locations on chromosomes, memes presumably exist in brains, and we have even less chance of seeing one than of seeing a gene (though, in an article referred to by Blackmore, the neurobiologist Juan Delius had pictured his conjecture of what a meme might look like). As with genes, we track memes through populations by their phenotypes. The 'phenotype' of the Chinese junk meme is made of paper. With the exception of 'extended phenotypes', such as beaver dams and caddis larva houses, the phenotypes of genes are normally parts of living bodies. Meme phenotypes seldom are.

But it can happen. To return to my school again, a Martian geneticist,

visiting the school during the morning cold bath ritual, would have unhesitatingly diagnosed an 'obvious' genetic polymorphism. About 50 per cent of the boys were circumcised and 50 per cent were not. The boys, incidentally, were highly conscious of the polymorphism and we classified ourselves into Roundheads versus Cavaliers (I have recently read of another school in which the boys even organised themselves into two football teams along the same lines). It is, of course, not a genetic but a memetic polymorphism. But the Martian's mistake is completely understandable; the morphological discontinuity is of exactly the kind that one normally expects to find produced by genes.

In England at that time, infant circumcision was a medical whim, and the Roundhead/Cavalier polymorphism at my school probably owed less to longitudinal transmission than to differing fashions in the various hospitals where we happened to have been born – horizontal memetic transmission, yet again. But through most of history circumcision has been longitudinally transmitted as a badge of religion (of *parents'* religion I hasten to point out, for the unfortunate child is normally too young to *know* his own religious mind). Where circumcision is religiously or traditionally based (the barbaric custom of female circumcision always is), the transmission will follow a longitudinal pattern of heredity, very similar to the pattern for true genetic transmission, and often persisting for many generations. Our Martian geneticist would have to work quite hard to discover that no genes are involved in the genesis of the roundhead phenotype.

The Martian geneticist's eyes would also pop out on stalks (assuming they were not on stalks to begin with) at the contemplation of certain styles of clothing and hairdressing, and their inheritance patterns. The black skull-capped phenotype shows a marked tendency towards longitudinal transmission from father to son (or it may be from maternal grandfather to grandson), and there is clear linkage to the rarer pigtail-plaited sideburn phenotype. Behavioural phenotypes such as genuflecting in front of crosses, and facing east to kneel five times per day, are inherited longitudinally too, and are in strongly negative linkage disequilibrium with each other and with the previously mentioned phenotypes, as is the red-dot-on-forehead phenotype, and the saffron robes/shaven head linkage group.

Genes are accurately copied and transmitted from body to body, but some are transmitted at greater frequency than others – by definition they are more successful. This is natural selection, and it is the explanation for most of what is interesting and remarkable about life. But is there a similar meme-based natural selection? Perhaps we can use the Internet

again to investigate natural selection among memes? As it happens, around the time the word 'meme' was coined (actually a little later), a rival synonym, 'culturgen', was proposed. Today, culturgen is mentioned twenty times on the World Wide Web, compared with memetic's 5042. Moreover, of those twenty, seventeen also mention the source of the word, falling foul of the *Oxford English Dictionary*'s criterion. Perhaps it is not too fanciful to imagine a Darwinian struggle between the two memes (or culturgens), and it is not totally silly to ask why one of them was so much more successful. Perhaps it is because meme is a monosyllable similar to gene, which therefore lends itself to quasi-genetic sub-coinings: meme pool (352), memotype (58), memeticist (163), memeoid (or memoid) (28), retrometne (14), population memetics (41), meme complex (494), memetic engineering (302) and metameme (71) are all listed in the 'Memetic Lexicon' at *http://www.lucifer.com/virus/memlex.html§MEME* (the numbers in parentheses count the mentions of each word on the Web on my sampling day). Culturgen-based equivalents would be more obvious but less snappy. Or the success of meme against culturgen may have been initially just a non-Darwinian matter of chance – memetic drift (85) – followed by a self-reinforcing positive feedback effect ('unto every one that hath shall be given, and he shall have abundance; but from him that hath not shall be taken away even that which he hath', Matthew 25: 29).

I have mentioned two favourite objections to the meme idea: memes have insufficient copying fidelity, and nobody really knows what a meme physically is. A third is the vexed question of how large a unit deserves the name 'meme'. Is the whole Roman Catholic Church one meme, or should we use the word for one constituent unit such as the idea of incense or transubstantiation? Or for something in between? Susan Blackmore gives due attention to such questions, but she rightly concentrates on a more constructive approach, developing the positive explanatory power of the 'memeplex' – an abbreviation which she prefers over the full 'coadapted meme complex', and I shall be surprised if in time her book does not bring about a Darwinian reversal of their numerical fortunes (today, 20 and 494, respectively).

Memes, like genes, are selected against the background of other memes in the meme pool. The result is that gangs of mutually compatible memes – coadapted meme complexes or memeplexes – are found cohabiting in individual brains. This is not because selection has chosen them as a group, but because each separate member of the group tends to be favoured when its environment happens to be dominated by the others. An exactly similar point can be made about genetic selection.

Every gene in a gene pool constitutes part of the environmental background against which the other genes are naturally selected, so it's no wonder natural selection favours genes that 'cooperate' in building those highly integrated and unified machines called organisms. Biologists are sharply divided into those for whom this logic is as clear as daylight, and those (even some very distinguished ones) who just do not understand it – who naïvely trot out the obvious cooperativeness of genes and unitariness of organisms as though they somehow counted against the 'selfish gene' view of evolution. Susan Blackmore not only understands it, she explains the matter with unusual clarity and goes on to apply the lesson with equal clarity and force to memes. By analogy with coadapted gene complexes, memes, selected against the background of each other, 'cooperate' in mutually supportive memeplexes – supportive within the memeplex but hostile to rival memeplexes. Religions may be the most convincing examples of memeplexes but they are by no means the only ones. Susan Blackmore's treatment is, as ever, provocative and revealing.

I believe a sufficient case has been made that the analogy between memes and genes is persuasive and that the obvious objections to it can be satisfactorily answered. But can the analogy do useful work? Can it lead us to powerful new theories that actually explain anything important? This is where Susan Blackmore really comes into her own. She warms us up with some fascinating vignettes which get us used to the memetic style of reasoning. Why do we talk so much? Why can't we stop thinking? Why do silly tunes buzz round our heads and torment us into insomnia? In every case she begins her response in the same way: 'Imagine a world full of brains, and far more memes than can possibly find homes. Which memes are more likely to find a safe home and get passed on again?' The answer comes back readily enough, and our understanding of ourselves is enriched. She pushes on, with patience and skill applying the same method to deeper and more exacting problems: What is language for? What attracts us to our mates? Why are we so good to each other? Did memes drive the rapid, massive, and peculiar evolutionary expansion of the human brain? Along the way, she shows how the theory of memes can throw light on particular areas where she has special expertise from her academic career as a psychologist and sceptical investigator of the paranormal: superstition and near-death experience.

In the end, showing greater courage and intellectual *chutzpah* than I have ever aspired to, she deploys her memetic forces in a brave – do not think foolhardy until you have read it – assault on the deepest questions of all: What is a self? What am I? Where am I? (famous questions posed

by Daniel Dennett long before he became the philosophical mentor of all meme theorists). What of consciousness, creativity and foresight?

I am occasionally accused of having backtracked on memes; of having lost heart, pulled in my horns, had second thoughts. The truth is that my first thoughts were more modest than some memeticists, including perhaps Dr Blackmore, might have wished. For me, the original mission of the meme was negative. The word was introduced at the end of a book which otherwise must have seemed entirely devoted to extolling the selfish gene as the be-all and end-all of evolution, the fundamental unit of selection, the entity in the hierarchy of life which all adaptations could be said to benefit. There was a risk that my readers would misunderstand the message as being *necessarily* about genes in the sense of DNA molecules. On the contrary, DNA was incidental. The real unit of natural selection was any kind of *replicator*, any unit of which copies are made, with occasional errors, and with some influence or power over their own probability of replication. The genetic natural selection identified by neo-Darwinism as the driving force of evolution on this planet was only a special case of a more general process that I came to dub 'Universal Darwinism'. Perhaps we would have to go to other planets in order to discover any other examples. But perhaps we did not have to go that far. Could it be that a new kind of Darwinian replicator was even now staring us in the face? This was where the meme came in.

I would have been content, then, if the meme had done its work of simply persuading my readers that the gene was only a special case: that its role in the play of Universal Darwinism could be filled by any entity in the universe answering to the definition of Replicator. The original didactic purpose of the meme was the negative one of cutting the selfish gene down to size. I became a little alarmed at the number of my readers who took the meme more positively as a theory of human culture in its own right – either to criticise it (unfairly, given my original modest intention) or to carry it far beyond the limits of what I then thought justified. This was why I may have seemed to backtrack.

But I was always open to the possibility that the meme might one day be developed into a proper hypothesis of the human mind, and I did not know how ambitious such a thesis might turn out to be. Any theory deserves to be given its best shot, and that is what Susan Blackmore has given the theory of the meme. I do not know whether she will be judged too ambitious in this enterprise, and I would even fear for her if I did not know her redoubtable qualities as a fighter. Redoubtable she is, and hard-nosed too, but at the same time her style is light and personable. Her thesis undermines our most cherished illusions (as she would see them)

of individual identity and personhood, yet she comes across as the kind of individual person you would wish to know. As one reader I am grateful for the courage, dedication and skill she has put into her difficult task of memetic engineering, and I am delighted to recommend her book.