

LONDON
SCHOOL of
HYGIENE
& TROPICAL
MEDICINE



LSHTM Research Online

Aunger, R; (2006) What's the matter with memes? In: Grafen, A; Ridley, M, (eds.) Richard Dawkins: How a Scientist Changed the Way We Think. Oxford University Press, Oxford. <https://researchonline.lshtm.ac.uk/id/eprint/11082>

Downloaded from: <http://researchonline.lshtm.ac.uk/11082/>

DOI:

Usage Guidelines:

Please refer to usage guidelines at <https://researchonline.lshtm.ac.uk/policies.html> or alternatively contact researchonline@lshtm.ac.uk.

Available under license: <http://creativecommons.org/licenses/by/2.5/>

<https://researchonline.lshtm.ac.uk>

What's the matter with memes?

Robert Aunger

LSHTM

I'm old enough to be one of the first generation of people brought to the 'gene's-eye-view' of biology by reading *The Selfish Gene*. It changed my outlook on life and has had a profound influence on my subsequent thinking. This transformative experience has been reproduced in many thousands of readers over the past 30 years. Oddly enough, for someone who has studied memes, I don't remember being particularly taken with the final chapter, 'The Long Reach of the Gene', at the time. What transported me then was the profound general outlook the book provided on why social interactions work the way they do. My professional concerns with the meme concept, introduced by Dawkins in that last chapter, came later, when I became interested in understanding cultural change.

A meme, of course, is defined as the fundamental unit of cultural transmission. From an evolutionary perspective, it plays the role in cultural change equivalent to that of the gene in biological change: as the basic unit of inheritance allowing the accumulation of adaptations. The idea is that, like a gene, a meme is a *replicator* (a concept also first defined by Richard Dawkins in *The Selfish Gene*). Genes replicate through the duplication of DNA strands; cultural replication, or the duplication of memes, takes place through the social transmission of information.

Dawkins was not the first scholar to broach the idea that culture might be underpinned by the replication of bits of information. The idea had been in the air for some time, with a variety of linguistic novelties being coined to describe a cultural replicator over the years: 'culturgen', 'mnemotype', 'culturetype' and 'sociogene'. But Dawkins' use of the word 'meme' caught on. Indeed the story of the spread of the meme meme makes a good case history in memetics (the study of memes).

To show the success of his invention, Dawkins (1999) did a web search in 1998 in which 'memetic' (used to eliminate possible confusion with the French word 'même') returned 5,000 web pages while 'culturgen' (the main contemporary rival, fashioned by Lumsden and Wilson in their book *Genes, Mind and Culture*) returned only 20. In the intervening seven years, the comparative advantage of meme has increased dramatically: memetic now appears on over 168,000 web-pages, while culturgen (or culturegen) lags with 537; other alternatives are nearly invisible. The 'meme' has won the contest to be the accepted name for the fundamental unit of culture. (A particularly convincing sign is the fact that E.O. Wilson uses 'meme' seven times in his book *Consilience*, while mentioning his own coinage, 'culturgen', only once.) The reason could be purely semantic: 'culturgen' is harder to say, while 'meme' easily blossoms into 'memplex' or 'meme pool'. On the other hand, perhaps the term's success is due to the fact that millions of people have now read *The Selfish Gene*.

The memetic banner has since been carried forward by a growing battalion. Since *The Selfish Gene* was first published, a number of books (by Blackmore, Dennett, Distin and myself) (Dennett 1991, 1995; Blackmore 1999; Aunger 2001, 2002; Distin 2004), numerous articles, an

electronic journal (the *Journal of Memetics*) and countless web postings and pages have been devoted to developing the meme meme.

However, when he introduced the idea, Dawkins' wasn't intending to inspire a new field of speculation and research; he was actually introducing an example of a second replicator, to show that Darwinian replication is not confined to genes alone (Dawkins 1999). He suggested that successful memes, like other replicators, should exhibit three crucial characteristics: fidelity, fecundity, and longevity (Dawkins 1976). Fidelity refers to the ability of a replicator to retain its information content as it passes from mind to mind. Fecundity is a measure of a replicator's power to induce copies of itself to be made. Longevity is less crucial; it only suggests that memes which survive longer have more opportunities to be copied, so the number of their offspring can increase too.

The meme meme seems to have these qualities in spades. Indeed, the meme idea has spread through both 'high brow' and popular culture, being used in a variety of ways by different disciplines or interest groups. For example, changes in the frequency with which the birds in an area sing elements of their song has been studied by animal behaviourists as a kind of 'population memetics' (by analogy to the study of changes in gene frequencies in population genetics). Similarly, computer scientists have argued that enabling one robot to imitate the behaviour another ('meme copying') is a way to get robots to develop 'culture'. Inspiring customers to spread good word-of-mouth about their products is also seen by some business writers as an exciting new tactic for increasing sales – a process they believe takes advantage of 'meme power'. Most infamously, perhaps, theologians have reacted to Dawkins' well-known atheistic stance and espousal of religious beliefs as harmful 'mind viruses' with defences of their beliefs – as in the book by John Bowker debating the question *Is God a Virus?*

Nevertheless, no significant body of empirical research has grown up around the meme concept (the bird-song work being the sole, limited exception), nor has memetics made empirically testable propositions or generated much in the way of novel experimental or observational data. In fact the memetic literature remains devoted almost exclusively to theoretical antagonisms, internecine battles, and scholastic elucidations of prior writings on memes. This is typically the sign of a science in search of a subject matter.

Why is memetic science ailing? I think most of the problems have to do with the lack of a useful definition. I would like to spend my time in this brief essay attempting to clarify this basic issue. As we will see, getting specific about the nature of memes leads to questions about whether there is indeed any subject matter for memetics to study.

So just what is a meme? Dawkins famously argued in *The Selfish Gene* that memes could be "tunes, catch-phrases, clothes fashions, ways of making pots or of building arches". This definition allowed memes to be found in various kinds of things: inside people's heads, in people's behaviour, and in artefacts. Susan Blackmore (2003) agrees with this broad definition in her book *The Meme Machine*, and argues that memes are implicated in the origins of human biology (particularly our large brains), culture (especially language, religion, art) and technology (which has become more and more efficient at copying and multiplying memes in artefacts like books and the World Wide Web). Basically, she sees memes as driving just about every interesting aspect of human evolution. This makes memes very powerful indeed. The problem is that, if memes explain everything, then they explain nothing. This sort of catch-all definition is too broad to be scientifically useful, and, I believe, accounts for memetics being empirically moribund at present.

I have argued, on the contrary, that what makes the meme concept special as an account of cultural evolution is its role as a *replicator* in culture (Aunger 2002). This is consistent with Dawkins' original objective in positing the existence of memes as a foil to genes.

The replicator concept has been one of Dawkins' lasting contributions to evolutionary theory. However, finding a way to define replication so that it encompasses all of the known replicators – genes, prions, computer viruses and memes – has been difficult. I have suggested (Aunger 2002) that replication can be defined as a special relationship between a source and a copy such that four conditions hold:

- causation (the source must play some role in bringing about the conditions that lead to a copy being made);
- similarity (the source and copy must resemble each other in relevant respects);
- information transfer (what makes the copy similar to the source must be derived from the source); and
- duplication (the source and copy must coexist for some time).

What does this definition of replication imply about the nature of memes? Does it restrict their definition in a useful way? I believe it does, but it will take a bit of analysis to see why.

Dawkins, like other memeticists, has argued that memes, like other replicators, can exist in many different forms. In effect, replicators are seen as symbolic entities which can morph from one form to another. Dawkins and others tell stories like the following, in which a gene is duplicated in a rather complicated fashion. Imagine a gene sequencing machine has 'decoded' a stretch of DNA into the familiar sequence of Gs,As,Ts and Cs (e.g., 'GCATACGATA'). This sequence is then printed onto a piece of paper, which is subsequently fed into another machine which reconstructs the same sequence of amino acids that made up the original DNA. That newly-created length of DNA is finally inserted into the nucleus of a cell and begins to function as evolution has designed it to.

In this example, the gene appears to have gone from being a portion of DNA to a sequence of markings on paper, then back to DNA. In effect, one code has been translated into another and back again, with the two different codes being realized in two different physical substrates. There is a one-to-one correspondence between the two coding systems, each of which has only four values, so high fidelity conversion back and forth is not difficult to achieve.

But let's look at this story more closely. There is certainly a causal chain in which information from one stretch of DNA is transferred to another stretch of DNA through the intermediate step of being stored symbolically on paper. Thinking of information in an abstract way suggests that the gene has been converted to a paper form, and that information inheritance has occurred: the crucial information seems to have been passed right down the line from one 'real' gene to another. However, it also appears nonsensical in evolutionary terms to argue that the symbol sequence on paper *is* a gene: the paper form does not conserve the essential features of a gene, its evolved functions (Hull and Wilkins 2001). In particular, the sequence of symbols can't produce a protein, or regulate the operation of other genes, no matter what environment the piece of paper is put in. That is because, in a different coding system on a different physical medium, this capability is lost. The symbol string *does*, however, hold information *about* a gene, which is used by specialized machinery as the basis for putting together the proper sequence of amino acids constituting that gene. So DNA → paper → DNA represents a causal

chain, but not an evolutionary lineage. This is because a lineage should constitute a sequence of copies, each of which is able to make further copies of things like itself – that is, working replicators, all along the line (Hull and Wilkins 2001).

How can we reconcile the fact that there is causation, information transfer and the duplication of DNA in this sequence but no evolutionary lineage? The only replication condition this example fails, according to our definition of replication above, is similarity: one copy must be 'like' the next. Is this lack of similarity due to the change in code from DNA to paper? Actually, there is a change in code during 'normal' DNA replication and expression: the duplication and transcription of DNA strands involve RNA (as primers or messengers, respectively); but RNA works via a slightly different coding scheme than DNA (changing one of the four nucleotides). There are also cases in which the replication of cultural information involves code-switching; for example, one can change codes of music from mp3 to wma or other forms while duplicating files. But all of these music codes are digital (two symbols only) and exist on the same medium: magnetic memory in a music player or computer.

So it isn't the change in code that matters to replication; it is the change in substrate. Replication appears to be substrate-specific (Aunger 2002). This is probably due to the fact that replication is a rather fragile process, a specialized kind of duplication which requires precise management – which means the beginning and end states must be physically similar, based in the same kind of substance. Certainly, no known replicator can replicate on more than one substrate: genes in DNA, prions as proteins, and computer viruses in computer memory. Presumably the same condition holds for memes, if they are replicators.

What is the proper substrate for memes then? It is commonly accepted that the primary repository of memes is brains. Why? Because memes are supposed to explain cultural change, and the quintessential cultural traits such as beliefs and values which distinguish one culture from another are in people's heads. Presumably for reasons like this, Dawkins, in his second book, *The Extended Phenotype*, restricted a meme to being 'a unit of information residing in the brain'.

But we still have a conceptual problem: if replicators are restricted to single substrates, how can we explain processes in which replicators appear to switch substrates, as in our story about a gene above? If genes don't exist in artefacts, then how can we account for a life history in which a gene passes through a phase in which it exists only as a piece of paper? How can a second copy of DNA acquire its genetic information (in our story above) if it has had no contact with its creator, the original bit of DNA? The answer is that the gene must be reconstructed from the information that *is* present in the symbolic sequence on paper, and which bears some relationship to the gene sequence. Fancy machines must reverse engineer the gene from the information on the piece of paper.

I have argued that a similar process occurs in the case of memes: when someone reads a book, and thereby acquires the author's ideas without ever meeting the author face-to-face, the book has served as a template, holding information which creates complex visual signals which, when perceived by the reader, instigates the reconstruction of the author's memes in the reader's mind (Aunger 2002). Just as the gene in the story above is reconstituted in DNA based on a paper-based representation, so too can a meme be reconstructed from a representation found on a piece of paper in a book.

Even face-to-face communication relies on the ability of human minds to engage in the reconstruction of information. This is because brains don't come into direct contact with one another. To jump the gap between minds, memes must use a signalling system, such as speech. This in turn means that message receivers must reconstruct a meme from the information contained in the signals it produces. The central question then is whether this mental reconstruction process can result in a copy of the original meme being produced – whether reconstruction satisfies the conditions for a replication process outlined above.

At present, it is difficult to know because we don't currently understand how social learning occurs. However, there are suggestions, largely from linguistics – the study of the most sophisticated natural signalling system known -- that 'copying the product' (as Blackmore calls it) is a process fraught with difficulty. The information contained in a message is rarely sufficient to establish its meaning. Each instance of listening to someone else requires inferring not only the semantic content of the message, but also the intentions of the speaker, which may bear little relationship to the message. For example, ironic communication is based on saying the reverse of what you mean (e.g., "I love your hairdo"). To make sure that the communication results in the receiver interpreting a message in something like the way the speaker intends, there have to be complex regularizing mechanisms (to eliminate spurious or extraneous elements) and a lot of shared background knowledge. So interpersonal communication is an instance of the same kind of process as learning from artefacts: a constructive process based on inadequate signals -- but this time received from an active, as opposed to an inert, interlocutor.

Even the most efficient form of social learning, imitation, which is supposed to ensure high-fidelity copying, is likely to introduce variation into what is learned (Sperber 2000). Exact copying is not a feature to be expected of everyday human communication because the signals we send are highly impoverished compared to what we infer from them. Certainly, many of the experimental studies of cultural transmission show rapid decay in messages, and reversion to 'lowest common denominator' content (Mesoudi 2005). If this is the case, it seems unlikely that culture can be viewed productively as the creation of lineages of information transmission with high-fidelity duplication and the long-term maintenance of cultural content. Human communication systems are thus unlikely to involve replicator-like inheritance – at least in the preponderance of cases. As a result, there may be no such thing as memes, in the strict sense.

More fundamentally, communication, when seen from an evolutionary point of view, is not *designed* to result in the copying of information. Another of Dawkins' major contributions was to point out that communication is a form of signalling designed to manipulate the minds, and hence the behaviour, of other animals (Dawkins and Krebs 1978). It is often in an individual's interest to get others to behave in ways which provide them benefits they can't achieve for themselves. This can be achieved by sending others information about a purported change in circumstance on which those individuals will then feel they should act. In some cases, what the communicator wishes won't be in the best interests of those listening, so the communicator will want to hide his or her true ambition -- not only from the message receivers, but often from themselves as well, so as to more 'honestly' signal their apparent, deceptive intent (Trivers 2000). From the message receiver's point of view, it will be important to make sure others are not trying to influence you in detrimental ways. Message receivers will only care about copying what is in someone else's head if that information is relevant to them, in their situation. But this won't often be the case, given that individuals are typically in different situations, with different interests (Sperber and Wilson 1995).

From this perspective, communication is not a peaceful exchange of information but rather tacit interpersonal warfare using information as a weapon. Of course when genetic or social interests overlap, communication can be cooperative, and information copying might be a desired outcome of message passing between cooperators. However, most cooperation requires people to adopt complementary, rather than similar, roles. Think of a simple example: two people trying to move a piano upstairs, one going backwards, the other forwards. Here, most of the shouting is about persuading the partner to move their bit of the piano to the left or right. Even in such cases, it doesn't seem necessary to know what is in the other fellow's mind to succeed. The knowledge of each cooperator can remain quite distinct.

Of course, even if communication isn't about replicating knowledge, memes might still be able to parasitize the communication process in order to duplicate themselves. However, if the objective of communication is primarily to manipulate those with different interests to oneself, natural selection should be expected to have evolved mechanisms for persuading others, not for copying information. So it might be difficult for memes to find ways to replicate information when duplication facilities have not evolved.

There is another difficulty to mention. We have relied throughout this discussion on the commonsensical assumption that memes occupy slots in the brain for different cultural traits. For example, the meme meme is one candidate for the 'name of unit of cultural transmission' slot, competing with culturgen as well as other terms as possible values. Folk psychology suggests the existence of such a concept, but perhaps the brain doesn't work that way; perhaps it represents information rather differently than the analogy to a filing system would imply. This possibility opens whole new vistas for what memes might be in conceptual terms: not units of language like words or even abstractions like concepts (such as the meme meme), but something our conscious minds cannot conceive of – perhaps something as alien to folk psychology as the computer representation of words in binary digits (Aunger 2002). The whole project of counting words on web pages or even instances of mental concepts in brains may be misguided. My own view is that memetics can only really take off once we have a better idea of how brains manage information, much like biology blossomed after the discovery of the DNA-based mechanism of gene replication. However, if it turns out that social learning typically doesn't involve the replication of information, then models of cultural evolution other than memetics will be necessary.

For a number of reasons, then, the replication of information is unlikely to be how most social learning occurs. Neither are memes necessary to explain cultural traditions. Henrich and Boyd (2002) have shown that even if copying is sloppy when individuals communicate with each other, the result of lots of sloppy social learning, when aggregated to the population level, can appear like a replication-based process in the sense that cultural traditions can still be maintained and adaptations can accumulate over time. This is true if one assumes that human psychology includes a tendency to favour the acquisition of specific trait values, or what Sperber calls 'cultural attractors'. Thus, even if memes aren't at work in culture, it can appear as if they were. So taking the stability of culture as *prima facie* evidence of the existence of memes is mistaken. Replication is not a necessary component of an interesting Darwinian process, and may not be involved in the explanation of human culture. Dawkins (1982a) presaged a similar conclusion long ago: "My own feeling is that its main value [the meme hypothesis] may lie not so much in helping us to understand human culture as in sharpening our perspective of genetic natural selection". My attempt to provide a more precise definition of memes has, ironically, shown that memetics appears to be in search of subject matter because its central claim, the meme hypothesis, lacks substance.

A final speculation about the fate of memes: even if it turns out that there are no mental replicators, it will be difficult to deny memes a role in the future of cultural evolutionary studies. This is because the meme meme has already become part of the culture it was supposed to explain – as attested by the frequency of its mention on the Web. I therefore suspect people will continue to use the word ‘meme’ in a vague way when discussing cultural change. But I also predict that memetics is unlikely ever to become an empirical science, because when we define memes in a manner precise enough to start making testable predictions, we find that we have largely defined them out of existence.

The last chapter of *The Selfish Gene* has thus proven incredibly provocative, and productive – at least in the sense of having spawned renewed interest, and a burgeoning literature, in the evolution of culture. At minimum, the meme concept has shown how evolutionary biology provides a model for the study of a central concept in the social sciences: culture. Interestingly, Dawkins (1982b) suggested that any process which showed design was likely to be due to the natural selection of random variants -- a principle he called ‘Universal Darwinism’. Some have taken this idea as a rallying cry, and used the meme concept as part of a general programme to apply Darwinian principles to the disciplines bordering on biology, particularly psychology and the social sciences. This kind of theoretical unification is highly desirable, if only for the parsimonious explanations it provides for a broad range of phenomena. But of course the idea that Darwinian theory can better account for the subject matter of a discipline than theories home-grown in that discipline itself has been fiercely resisted as a form of territorial imperialism by those whose territory is being contested (e.g., Sahlins 1976; Kitcher 1985). Nevertheless, the success of evolutionary psychology and cultural evolutionism are clear indications of the rapid spread of what might be called the ‘Universal Darwinian programme’ – and testimony to the fertile theoretical mind of Richard Dawkins.

REFERENCES

- Aunger, Robert (2002) *The Electric Meme*. New York: Simon and Schuster.
- Aunger, Robert, ed. (2001) *Darwinizing Culture: The Status of Memetics as a Science*. Oxford: Oxford University Press.
- Blackmore, Susan (2003) ‘The evolution of meme machines’. In Meneghetti, A, et al. (eds.), *Ontopsychology and Memetics*, Rome: Psicologica Editrice, pp. 233-40.
- Blackmore, Susan (1999) *The Meme Machine*. Oxford: Oxford University Press.
- Bowker, John (19???) *Is God a Virus?* PUBLISHER.
- Dawkins, Richard (1999) Introduction to Blackmore, 1999.
- Dawkins, Richard (1982a) *The Extended Phenotype*. Oxford: Oxford University Press.
- Dawkins, Richard (1982b) ‘Universal Darwinism’. In D.S.Bendall (ed.) *Evolution from Molecules to Men*. Cambridge: Cambridge University Press, pp. 403-25.
- Dawkins, Richard (1976) *The Selfish Gene*. Oxford: Oxford University Press.
- Dennett, Daniel C. (1991) *Consciousness Explained*. New York: Little, Brown.
- Dennett, Daniel C. (1995) *Darwin’s Dangerous Idea*. New York: Simon and Schuster.
- Henrich, Joseph and Robert Boyd (2002) ‘On modeling cognition and culture: Why replicators are not necessary to cultural evolution’. *Journal of Cognition and Culture* 2:87-112.
- Hull, David and John S. Wilkins (2001) ‘Replication’. *Stanford Encyclopedia of Philosophy*. <http://plato.stanford.edu/entries/replication/>
- Kitcher, Philip (1985) *Vaulting Ambition: Sociobiology and the Quest for Human Nature*. Cambridge, MA: MIT Press.

- Lumsden, Charles J., and Edward O. Wilson (1981) *Genes, Mind and Culture: The Coevolutionary Process*. Cambridge: Harvard University Press.
- Mesoudi, Alex (2005) *The Transmission and Evolution of Human Culture*. PhD thesis, University of St. Andrews.
- Sahlins, Marshall (1976) *Culture and Practical Reason*. Chicago: University of Chicago Press.
- Sperber, Dan (2000) 'An objection to the memetic approach to culture'. In Aunger, Robert, (ed.) *Darwinizing Culture: The Status of Memetics as a Science*. Oxford: Oxford University Press, pp. 163-174.
- Sperber, Dan, and Deidre Wilson (1995) *Relevance: Communication and Cognition*, 2nd ed. Oxford: Blackwell.
- Trivers, Robert (2000) 'The elements of a scientific theory of self-deception'. *Annals of the New York Academy of Sciences* 907:114-131.