

CHAPTER 11:

Conclusion

ROBERT AUNGER

This book seeks to determine whether the idea of memes might provide the foundation for a progressive line of research on cultural diversification and evolution. In this, the concluding chapter, I don't intend to make the definitive statement concerning the future of memetics. Rather, I attempt to make sense of what has gone before, and to find where there might be grounds for coming down on one side or another of the key issues identified by the authors of the preceding chapters. My comments will be arranged by academic discipline, as it is from these varying perspectives that problems naturally come into view. Indeed, the entire book is arranged -- by default -- in a similar fashion. It works out (roughly speaking) that meme promoters, put first in this book, tend to be biological in background or inclination, while the more critical voices dominating the later chapters tend to come from psychology and especially the social sciences. I follow the same order in setting out my comments here.

Evolutionary theory

Given its origin in the work of the zoologist and evolutionary theorist Richard Dawkins, the memetics literature has continued to exhibit the strong influence of evolutionary biology. Many of the problems with pursuing this line of research therefore arise from the analogy made between genes as the biological replicator and memes as its cultural equivalent.

Explaining cultural similarity

Just because the "meme meme" has been highly successful in popular culture and has even appeared in the Oxford English Dictionary, it has not yet been established whether memes are a subject worthy of scientific study. The approval of journalists and British lexicographers merely reflects common usage, and establishes memes as a viable folk psychological concept. However, we cannot be certain that memes -- considered as a scientific concept -- exist.

Why is this? Let's step back a second and take a look at what distinguishes memetics from alternative theories. Memetics asserts that we can take a "meme's-eye-view" with respect to the diffusion of culture. The obvious implication is that there is a previously unnoticed agent participating in social communication -- something besides just the sender and receiver that needs to be accounted for. In effect, memetics postulates the existence of an evolutionary agent -- a replicator -- that evolves in accordance with its own interests (which may be independent from those of either the sender or receiver of messages). Most would identify this agent as the message itself. So a meme must be thought of as a replicator which is active during social communication in such a way that it can influence its own reproduction. The problem is that no one has yet identified bits of information with these qualities.

Why posit the existence of such a thing? Because the fact of cultural similarity needs explaining. Everyone has had the experience of someone else expressing opinions similar to their own or behaving like they do. This suggests there are multiple copies of the information underlying that belief or behavior in the population. But how did this commonality arise? Was the relevant information transmitted to them by others? Or perhaps similar environments caused commonly-held information -- information placed in people's heads at birth by genetic inheritance -- to be expressed by anyone in that situation. Or perhaps each individual learned the relevant piece of information through earlier experiences with their natural surroundings, without having communicated with anyone or possessing that knowledge innately. In effect, there are three standard explanations for cultural similarity:

- transmission (cultural evolution through social learning)
- genes (biological evolution), and
- individual learning (which is convergent evolution through the analogue of mutation from a cultural perspective).

Memetics is associated with the first of these. So what we require to prefer the memeticists' explanation of cultural similarity is proof that cultures evolve thanks to the non-genetic inheritance of information. The problem, then, is eliminating the other mechanisms (just outlined) which might underlie the regeneration of cultural traits over time, but which do not involve a cultural replicator -- or, indeed, social learning of any kind. How can we discriminate between these alternatives?

The more radical evolutionary psychologists (e.g., Tooby and Cosmides 1992) favor the gene option. They would minimize the role of transmission altogether, emphasizing instead the stimulation of innate mental content by potentially simple ecological stimuli. In essence, they believe "cultural" traits are already in the brain, with only an environmental spark required for them to be expressed. What remains to be explained from the evolutionary psychological perspective is not social transmission dynamics, but recall dynamics: what kinds of responses do different environments cause to arise? Boyd and Richerson (this volume) disposed of this possibility by arguing that the corpus of human knowledge accumulates too rapidly to be purely genetic in origin. So it seems unlikely that genes -- through the instrument of the adapted mind -- will single-handedly account for culture. Boyd and Richerson (this volume) also claimed that individual learning in similar environments is an inadequate explanation of cultural similarity, because groups living in the same environment display different suites of cultural traits. This seems, therefore, to leave us with only the cultural transmission explanation -- in effect, memes must be invoked to explain cultural similarity. So what's all the fuss about?

In fact, there is another possibility (which is not standard, so I didn't include it above): niche construction and ecological inheritance (see Laland and Odling-Smee, this volume). Cultural groups living side-by-side may not be living in the same effective environment, because they have modified their natural surroundings in distinct ways. In this case, people learn their cultural traits through interactions with artefacts, rather than other people. Cultural groups living in the same environment differ, in this view, not because they learn the beliefs and values which distinguish them from each other, but because they are influenced by artefacts inherited from previous generations. These can even be types of artefacts which do not communicate information from ancestors to the present-day inhabitants of those environments (like books do). Instead, they could take the form of tools and the "built" environment, which only indirectly influence attitudes and beliefs. So by invoking our ability to manipulate the environment over the long term (an ability we share with many other species), we can continue to discount the role of memes in explaining culture acquisition -- even in the face of rapid technological improvements such as surround us today. The feed-forward effects of ecological inheritance, coupled with big, evolved brains able to manipulate the information stored by the activity of previous generations in the environment, can in principle explain the similarity within, and differences between, cultural groups.¹

Indirect evidence of memes?

In the face of these competing schools of thought, each with vocal and sophisticated adherents, I suggest that for the memetic hypothesis to be favored, we require evidence of some kind that memes exist. This evidence can either be direct or indirect. From indirect evidence, one can infer the existence of memes from the marks their activity leaves behind in the world; direct evidence would show us where memes are and what they look like.

Good indirect evidence for memes would consist of establishing that there is an independent dynamic to cultural change which cannot be assigned to the goal-directed activity of people. One would need to observe a directionality to cultural change which reflects the interests of a replicator battling with genes for control over human behavior -- memes. This is why memes are commonly invoked to explain maladaptive cultural traits, why advocates often gravitate toward examples of memes which seem "irrational" for individuals (like celibacy), and why memes get equated with viruses, to prompt the implication of induced morbidity in "hosts." The problem is that -- except for the odd trait here and there - - culture is overwhelmingly adaptive for people, allowing our species to dominate our home planet in fairly spectacular fashion. If memes are parasites, they must be symbiotic ones.

So, if memes exist, it is more likely that the course of human evolution should reflect the increasingly interdependent interests of genes and memes. An increasingly effective mutualism between these replicators should result in the human species becoming able to explore new ecological niches, thanks to the additional functionality granted to humans by their relationship to the memetic symbiont. In effect, contemporary humans should be exploiting a broader evolutionary "design space" -- or range of life-ways in which they can thrive -- than was possible before memes came along.

The obvious manifestation of synergistic niche expansion through gene-meme cooperation is the rapid increase in technological improvements associated with civilization. Indeed, this is what many would say is the best kind of indirect evidence for the operation of memes -- their (unspecified) role in artefact development (e.g., Gabora 1997).² Boyd and Richerson's examples (this volume) of incremental improvements in tools such as compasses would seem to prove that a series of artefactual forms can exhibit descent with modification -- or the passing of information through a chain of exemplars, forming lineages of information transfer and duplication. After all, such intricate implements show evidence of design, or adaptation to particular functions.

But there are still two kinds of explanations for such obvious design. Does it arise because the best-performing tools are artificially selected by people to reflect their own needs? Or alternatively, is their design the "natural" outcome of independent replicators (memes again) working to achieve a higher probability of replication -- primarily by becoming more useful to people? In other words, Does the evolution of technology reflect the will of people or the interests of symbiotic memes? It should be apparent that it would be very difficult to tease such closely twinned hypotheses apart.

Nevertheless, some memeticists identify artefacts as memes (e.g., Blackmore 1999; Conte, this volume; Sperber, this volume). Do artefacts fill the bill? As Sperber (this volume) has forcefully argued, three criteria are required for replication -- causal efficacy, similarity, and inheritance. Sperber explicitly means to exclude cases of reconstruction from memetics through the criterion of inheritance. Inheritance here means that the information leading to the copy being produced must be acquired from the source, rather than having the recipient reconstruct the requisite information for itself. This criterion must hold true whether memes are defined as being in the head or in the form of artefacts.

In fact, there seems to be an impressive array of mechanical replicators out there -- chain letters, photocopies, FAXes -- which meet all of Sperber's criteria for replication. Let's take the case of web-page downloads. Lots of web-pages get visited by Internet "surfers," but only in a few instances is the information found there downloaded to the surfer's local hard disk. Presumably, it is some aspect of the content of those pages which triggers the download -- and hence the replication of that content. Software ensures that the copying process, based on source information, is high-fidelity.

From the "artefact's-eye-view," this is replication, with people placed solely in the role of catalysts for the process. Photocopying is perhaps a more straight-forward example: ink-on-paper serves as a template for the copy; there is no phenotypic conversion involved. Indeed, the copying process is just like meiosis: direct replication of the memetic genotype, ink-network to ink-network. Whether anyone reads the copy is immaterial. People are only needed to push the copy button on the machine. (Computer viruses replicate through networked computers with even less human involvement.) The important thing is that there are more copies of the artefact around at the end of the exercise. No replication of information

need occur in brains during the process, since each push of the copy button may be produced by some previously-installed mental rule about what to find appealing on the piece of paper.

To see artefacts as replicators, we must make a mental flip of perspective, to see the world as the replicator sees it. Dawkins (1976) taught us that we must often think of the biological process from the genetic replicator's point of view (which sometimes means that individual organisms become almost invisible -- as in the case of oncogenes, which reproduce themselves through a renegade cell lineage expanding at some cost to the individual's health). So too, from the replicating artefact's point of view, these substances -- rather than being external stores of information for the use of people, or aids to getting memes from one brain to another -- become the focus of a replication story. What is crucial to photocopying, for example, is the original on the glass, the spinning electrostatic drum, and the push of the copy button. The massive human brain is relegated in this story to the trivial task of button-pusher (which a simple-minded automaton could also do). Artefacts can in fact inherit information from other artefacts, and "a scholar is just a library's way of making another library" (Dennett 1995:346).

Memeticists (like Blackmore 1999, this volume) have invoked mental memes to explain the evolution of technology like photocopiers, but now we have just the reverse: the suggestion that technological replicators are being produced without a necessary role for mental replicators. The "best" evidence for memes -- the evolution of modern technology -- turns out to be an instance of replication from another point of view, and hence cannot be used to support the hypothesis of brain-to-brain replication. Rather, technological replication falls into the category of niche construction. Laland and Odling-Smee's ecological inheritance (this volume) may occur through such instances of artefact replication (although it can also occur through the mere survival of existing artefacts, if they last longer than a human generation). And the evolution of the large human brain -- a conundrum supposedly explained by memes inciting the construction of a bigger home for themselves, according to Blackmore (1999, this volume) -- turns out to be unnecessary for that process to occur, and so must be due to some other cause.

It thus appears that replication is happening all over the place -- inside cells (genes), between proteins (prions), and in the environment (artefactual replication). The irony is that it may not be happening in the way originally envisioned by Dawkins -- through social learning. Mind-to-mind replication may in fact be the least likely mechanism for replication (see the section on memetic phenotypes below). Indeed, whether memes exist in minds remains an open question. Certainly, no model in the memetic literature which makes a brain the site of replication meets Sperber's criteria.³

A terminological question now crops up: Should we call this technological replication of patterns on paper or in hard disks a memetic process? Certainly, the information in these patterns does not replicate via imitation, even broadly conceived, and therefore does not fit the original Dawkins/Blackmore definition. If culture is composed of information in people's heads, then the duplication of artefacts doesn't necessarily help us to explain culture. People may learn from these artefacts or not. Because the present book concerns memes as a contender for the explanation of cultural evolution, I will restrict my use of the word meme to information replicated through social learning (its original context), and leave the question of what to call the technological replication of artefacts to others.

Direct evidence of memes?

So our search for indirect evidence of memes actually led to the discovery of artefactual replicators and grim forebodings about the need for, or existence of, brain-based memes. We come down, I think, to the need for direct evidence of memes-in-minds to prefer the meme hypothesis for explaining culture. Since memes are replicators, they must be defined essentially by their means of replication, which should be distinct and independent from those of other replicators (including artefacts). Thus, in my view, the case for memes cannot be made without reference to a mechanism by which information is faithfully replicated through social transmission.

What does “mechanism of replication” mean in this context? By definition, it is the means by which information exerts some influence over the probability of it being reproduced (Dawkins 1982:83). One could go further and require a specification of the various resources and their roles in the replication process -- the steps leading to the product being assembled, and their speed -- but that is no doubt a task for the future.

I therefore conclude that only by finding a mechanism of replication which generates the similarity between people's beliefs and values can instances of inheritance-through-transmission be conclusively distinguished from something like the genetic or developmentalist (evolutionary psychological) alternative. This makes Blackmore's “existence proof” for memes, as presented in her contribution, unacceptable. It is based simply on the dictionary definition of memes, with a note that this definition implies memes are replicators. In fact, the involvement of memes in the maintenance and diffusion of mental culture remains an open question.

So I submit that meme-promoters will only be proven right about cultural inheritance when someone finds a meme. Nothing except seeing identifiable memes in action is likely to convince people sitting on the other side of the fence that memetics is the best option.

I also think it will be difficult to find a meme without specifying what the search is looking for, and where. Hull (this volume) says we don't need to have a crystal clear definition of memes in order to work with them. He (and Blackmore 1999:56) cite the oft-mentioned parallel example of genes: that purely operational definitions of genes during the first part of the twentieth century were sufficient for good science to be done. Unlocalized, metaphorical units of inheritance were certainly enough for Darwin to sweep all contenders aside in the nineteenth century, given the logical force of his argument for natural selection as a mechanism. So Hull's admonition to contemporary would-be memeticists is simply to go out and collect evidence of memetic activity in the social world.

Is that going to be good enough? I suggest not. In my view, the situation with respect to cultural inheritance is not the same as that for genes because genes are already established as a mechanism for informational inheritance. Once genes are on the scene, all inheritance, including cultural, might already be accounted for (although I agree with Boyd and Richerson that this is unlikely). If not, then we still have the option of invoking ecological inheritance. So identifying a more-than-operational meme and its mechanism of replication are both necessary before memetics can get off the ground. Only by providing a physical model of meme replication can memetics take its rightful position in the list of replicators covered by what Hull terms “general selection theory.” Until then, they remain simply an analogy to the better-known case of genes.⁴

Memetic phenotypes and the communication problem

Even if we ignore these empirical difficulties, major problems remain in meme theory. One is establishing how the genotype/phenotype distinction might work for memes. This distinction is crucial because brains don't directly infect each other with bits of brain-stuff; rather, they use signals or messages instead. Brain-to-brain transmission therefore necessarily involves the translation of memetic information from brain language to signal language, from one form or code to another, and back again. I will call this the “communication problem.”

There is also another reason memetics should concern itself with establishing what a memetic phenotype might be. The functional distinction between genotype/phenotype in the genetic system has been generalized by Dawkins and Hull as the replicator/interactor distinction (see my Introduction). Although it is possible for a replicator to also serve as an interactor (as ribosomes do, for example), such a situation is generally considered unlikely to persist. This is because replicators and interactors have fundamentally different roles to play in the evolutionary drama (as store of information and as survivor/transmitter, respectively), and it is usually inefficient for the same entity to play both roles. So a competitor system with independent replicators and interactors would almost certainly win out in an

evolutionary race, if only because a more specialized replicator would likely be more robust in its ability to duplicate itself. If memes are considered well-developed replicators, then memeticists will have to develop a notion of a memetic interactor, or “phenotype” (by analogy to the biological phenotype). While there are a number of contenders for this role, none has achieved widespread recognition.

Part of the problem with developing a rigorous notion of a memetic interactor is coming up with a criterion which surely identifies it as distinct from its progenitor, the memetic replicator. David Hull (this volume) put forward one criterion for making the distinction between a replicator and its interactor which is generalizable regardless of substrate (and thus a candidate for Universal Darwinism): the relative difficulty of reconstituting the replicator from an interactor. This is a generalization of the Weismannian notion that, in informational terms, you can't go “backwards” from protein to gene. Such an inability arises because there tends to be some slop in the production of phenotypes: genes don't code for one phenotype, they code for a gradient of possible variant forms (what biologists call a reaction norm), thanks to the impact of environmental conditions on development. So the relationship between replicators and their products is not one-to-one. This implies that information will be lost in the translation from meme to phenotype. It is this loss of information which makes the project of “reverse engineering” (or inferring the assembly instructions from seeing the product, as Susan Blackmore puts it) so difficult.

Clearing up what is a cultural replicator and what is a cultural interactor will also go some way toward avoiding the perennial confusion surrounding “Lamarckianism” in cultural evolution. Since the Lamarckian principle involves the inheritance of phenotypic variation, determining whether cultural evolution is Lamarckian depends on distinguishing between memotype and phenotype. Memes may change code or form during transmission, but cultural inheritance will be Lamarckian only if the meme is in phenotypic (informationally-compromised) form during transmission. In this case, the meme-recipient will acquire a phenotypic variant. So making the proper distinction between replicator and interactor forms is crucial for basic understanding in memetics.

However, this leaves us in an unfortunate quandary -- at least so long as we use information loss as the criterion for identifying the phenotypic form of a replicator. This quandary arises because, as Hull (this volume) notes, without a clear idea of what memetic information is -- that is, how the information in a bit of writing differs from the information in the piece of paper on which it is written -- we don't have a good way of determining when it is being lost. If we insist on using information loss as the defining criterion of interactors, progress in memetics will be inhibited until we know how this loss occurs.

Dan Sperber has argued that it is hard for a replicator to solve the problem of information loss during social communication. Artefactual replicators in the form of ink on paper can duplicate with very high fidelity: using photocopiers, we have direct replicator-replicator reproduction, and consequently no loss of information. However, as noted earlier, since bits of brain don't themselves make the journey from one head to another, the memetic life cycle requires that memes be translated from some neuronal construct into another form for social transmission -- for example, into parts of speech. Thus, memetic replication cycles involve stages of translation from one code and substrate to another. Since translation is rarely perfect, this implies that information leakage should regularly occur.

The problem with this is that, if speech is a phenotype, then it is compromised as a message carrier (this is the famous Chomskian “poverty of the stimulus” argument concerning linguistic message-passing). But then, for the sender's intent to be properly communicated, the receiver must compensate for this information loss by engaging in some kind of reconstitution of the message's intended meaning. However, if there is significant reconstruction of the information content of a meme by each host brain, then the likelihood of message replication is low, thanks to the vagaries of how each brain processes incoming information (due to the different background information individuals have acquired, the inferencing algorithms they use, etc.).

One way out of this problem, suggested by Sperber, would be for the brain to have a general decoder -- a utility enabling it to reliably infer the intention of the sender, and hence the substance of the message, regardless of any intervening noise during transmission or idiosyncracies of sender coding and production. In this view, brains should have evolved filters to assess the utility of information coming in

from the social environment to keep us from rapidly being swamped with bad information (or duped into stupid behaviors by people with ulterior motives). This normalizing inferential machine might also ensure the replication of memetic material during social transmission. However, its operation is unlikely to be perfect, so a high mutation rate remains a potential problem.

The need to communicate memes between brains through intermediaries also introduces another, more fundamental, complication. If psychological normalization of memetic inputs is important for successful communication, then memetic information is not, strictly speaking, inherited because it is not passed from person A to person B. Instead, the similarity of socially-acquired information between individuals has another cause: inherently structured inferential processing by the brain (Sperber, this volume). These reconstructive processes depend on a long history of genetic selection on the human cortex, not the passing of information from person to person in cultural lines of descent. In effect, the cause of the similarity between the information in A's and B's brains is the result of evolutionary psychology, not memetics. Since the causes are different, one can expect the population-level dynamics to also be different, thanks to differing rates of mutation or types of selection, for example. This creates a fundamental problem for memetics as an inheritance process (the general view of memetics).

However, the memetic process -- even if dependent on error-correction routines in the brain to produce the cultural similarity of beliefs and values -- still confers an evolutionary advantage. This is because the same information is acquired through transmission-plus-correction more efficiently and cost effectively than individual learning through trial-and-error could have done (Dan Dennett, personal communication.). Further, error correction is an important aspect of genetic inheritance as well, so replicator systems can operate with such assistance without having to be called something else.

Susan Blackmore (personal communication) notes that Sperber's reasoning leads to the expectation that, if there is a cultural replicator, there should also be selection for improved mechanisms for its transmission over time. In this way, the reliance on reconstituting information from local resources each iteration would be reduced and the proportion of information actually being transmitted increased. Her presumption is that this is indeed what has happened during the major transitions in cultural evolution, such as language, writing, and computer-based communication. But whether these have increased the transmissibility of memes, or merely their copying fidelity, remains to be determined.

Psychology

Another major set of issues concerns the psychology of memes.

Must we go inside?

A fundamental question in this domain is whether memetics can proceed without a clear idea of what kinds of transformations memes might undergo during storage and retrieval by brains. Can memetics leave the brain as a black box, and deal only with social transmission aspects? The virtue of ignoring psychology is that we need not worry about something we don't know too much about anyway: how the brain processes information. This is the line taken by Blackmore and Hull (this volume), who argue that memetics can cheerfully ignore what is going on inside people's heads because the real action is happening in the social sphere, or at the level of the population. Boyd and Richerson (this volume), sensing difficulties in this area for memetics -- that the psychological mechanisms underlying inheritance are likely to be messy and remain largely unknown -- shy away from these particulars. They claim that however the psychological side of things plays out, cultural evolution can nevertheless be seen as a Darwinian process from the level of the population: each generation somehow has to cause information to get stored in the brains of subsequent generations. And it is true that whatever is happening "inside" can be glossed as some kind of decision-making bias favoring one variant over another during transmission (which is effectively what gene-culture coevolutionary models do). But this does not very effectively limit the kinds of models which need to be investigated.

Further, if memetics disregards psychology, and there are major transformative processes at work in the brain, then memetics is only explaining part of the cultural evolutionary process. Since the survival of a meme might depend on an interaction between what happens to it inside and outside the brain, by ignoring one half of the picture memeticists may get the part they explicitly deal with -- the public or social part -- empirically wrong. Psychologically-oriented memeticists generally feel that no social theory, including memetics, can succeed without a proper psychological underpinning.

So, if we agree that we must have a mechanism producing similarity (as I argued above), then we can answer the question of whether memetics must involve itself in psychological issues. The answer is yes: it is crucial that we learn how we learn to become culturally competent members of society. Conte, Sperber and Plotkin are right, in my view, in this respect. I thus conclude that memetics must peek inside people's heads. Score a point for the psychologists.

Unfortunately, psychologically "realistic" population-level models of the cultural evolutionary process -- whether analytic, or the sort of computer-based simulations preferred by Conte -- remain for the future. This is because few social psychologists are interested in filling in the picture with regard to transmission biases, so the wait for greater psychological realism may be a long one.

Imitation

Related to the issue of how memes might replicate is the relationship of memetics to imitation. Two interconnected questions pop up here. First, Is a complicated brain essential for imitation? This issue is important because it determines who gets to have memes: only complex intentional agents like people, or more lowly creatures without cortices, such as birds? Many (including Plotkin) argue that there is no consensus concerning the psychological mechanisms of imitation. This is significant because, as Conte (this volume) says, you cannot define imitation without reference to the mental abilities involved. Using behavior as the sole criterion leads to confounds. For example, automatic contagion (such as yawning when others do) is direct phenotypic copying without the inferencing of mental contents. Counting contagion as a kind of imitation suggests that agents don't need to correctly infer another's intention (plus her beliefs and needs, etc.) in order to adopt or imitate her behaviors. What psychological resources imitation demands remains unknown.

The second aspect of the imitation question is a point on which many here voiced an opinion, so it appears to be central: Should memetic transmission be restricted to imitation? Blackmore, citing Dawkins as an authority, restricts memetics to cases of imitative behavior because, she asserts, only imitation serves as a direct copying process, and if memetics is to be founded on replication events, then only imitation can be counted as a memetic mechanism. But as we have just seen, the jury is still out on whether imitation is behavior copying or mental state inferencing (as assumed in the "theory of mind" literature). This leaves Blackmore's contention somewhat up in the air. Partly on these grounds, Boyd and Richerson, Conte, Hull, Laland and Odling-Smee, and Plotkin make attacks -- at least in passing -- at Blackmore's position on this issue. It seems that numbers, at least, are against her in this respect. The counter-proposals of Laland and Odling-Smee, Plotkin and Conte are particularly compelling, coming as they do from within the psychological fraternity.

Thus, there is little support for Blackmore's contention that memetics should be limited to imitation because imitation is the only mechanism which can support good replication. It may turn out that directly modelling the behavior of others is not more efficient than independent learning based on environmental cues. Basing meme replication -- "by definition" -- on imitation, as Blackmore does, is just not going to work. Imitation is too vague a gloss for what happens during (some kinds of) social transmission. At present, the process appears to involve a magical elision of mental substance from one brain to another -- much like the sympathetic transference or "contagion at a distance" characteristic of "primitive thought," according to some anthropologists (Hallpike 1979). Once the black box of imitation is opened up, we may find the magic disappears, and rather mundane mechanisms are at work.

Given this general discontent, it seems that any form of social learning, rather than imitation alone, is a better psychological foundation for the cultural evolutionary process. Reader and Laland (1999) take the famous example of milk bottle-top opening by birds as evidence of the need for this generalization. The pecking of bottle-tops has now gone on for many bird-generations, and spread through several European countries. Since it is generally felt that birds learn this bit of cleverness not by observing others, but by seeing opened bottle-tops, which inspires their own creativity (a process psychologists call "stimulus enhancement"), it seemed a pity to exclude such an example from the purview of memetics by limiting it to imitation-based diffusion.

However, if this liberal position on social learning is adopted, many repercussions ensue. For example, the phylogenetic history of memes suddenly becomes considerably longer, with birds and perhaps even more "primitive" creatures being allowed to have meme-based "protocultures." In addition, it means that direct contact between hosts is no longer required for memetic transmission, since the source of a meme (such as the tit which pecked a bottle-top) can be absent when a new, naive tit arrives on the doorstep. It is the artefact left behind -- that is, the pecked bottle-top itself -- which serves as the proximate stimulus for transmission of the pecking meme to the new arrival.

Allowing memes to be learned through any social mechanism implies, then, that memetics must address the issue of artefact production, since memes can be associated with these constructions, and not just brains. Laland and Odling-Smee (this volume) hinted at the importance of artefacts with their concept of niche construction. I discussed their concept from the "artefacts'-eye-view" previously. In artefact replication, humans are catalysts -- they push the "start" button which sets the process in motion. But now we see that memes can interact with artefacts as well, in their efforts to find new hosts. I suggest that the involvement of artefacts in a meme's life cycle can be seen as an elaboration of a more primitive process of memetic replication through signalling. In memetic social communication, a human source produces the catalyst -- a signal such as a gesture or bit of speech -- which causes a meme to be replicated in another brain. Such signals are not memes themselves, but rather moving memetic enzymes, produced by memetic activity in the message-sender's brain. On encountering the proper conditions -- to wit, an "innocent" brain -- this in-coming message instigates the meme replication process in the new host.

This simple model of memetic replication through communication becomes more complex when what might be called a "communicative artefact" steps into the middle of the communication process. In this case, message-senders create artefacts rather than signals -- for example, written messages rather than speech. These artefacts lay "dormant" in the environment, during which time they lie in wait for new hosts to infect. For example, words printed on paper can serve as a template for ambient light striking the paper, creating a catalytic signal that passes from the paper-artefact to a naive individual's eyes. This recipient individual, in turn -- and in good Sperberian style -- reconstructs the meme based on this "impoverished stimulus," using local mental resources. In this way, memes don't need to physically pass from brain to brain, and so don't need to adopt phenotypic forms such as signals themselves. Nevertheless, the "replication through communication" problem is solved because a new meme appears in the recipient brain which is causally connected to the source meme through the information provided by the message-catalyst. Sperber's inheritance condition is thus satisfied. And the process is replication as defined by Dawkins (1982; cited earlier) because the message has influenced the likelihood of a meme copy appearing; indeed, that is exactly the role a catalyst should perform in such a chain of events.

But remember that words printed on paper can also be part of a chain of artefact replication, as when they get photocopied. Thus communicative artefacts are the junction-point for two replication processes: that reproducing the artefacts themselves, and that producing new meme-copies. So technological development -- at least in cases of communicative artefacts -- can reflect the evolutionary interests not only of the artefacts themselves, but also those of the memes and people that interact with them. The difficulties of dealing with meme-artefact interactions are rarely discussed in the memetic literature, but such a complex phenomenon obviously requires attention if a comprehensive picture of meme replication is to be achieved.

Mental Darwinism

A second major point of contention is whether the memetic dynamic can be extended into the brain. Can we call individual learning a selection process just like the social transmission process (Changeux 1997 [1985])? This proposal has met with considerable disdain, and -- at least among academic psychologists -- is definitely a minority position (Henry Plotkin, personal communication). Blackmore, in particular, is adamant that whatever is happening inside the head should not be considered part of the memetic process; even if decision-making is in fact selectionist, it still should be treated as an independent replicator system in her view. This may be wise, given the possibility that the meme concept becomes vacuous when extended to cover replication in too many contexts.

However, including selection among alternative mental representations as an intrinsic part of the life cycle of a meme may be crucial to a successful memetics. Two benefits result from this conceptual move. First, only through an analysis of mental properties and processes can good models of transmission mechanisms, such as imitation, be understood. Second, by extending the Darwinian process into the brain ("Mental Darwinism"), one can avoid the confusion of thinking being called "directed," "intentionalist" or "Lamarckian." Instead of invoking a wholly new kind of process, one can simply suggest that decision-making is selection among memetic alternatives tossed up by some variation-producing process. In the end, the same substrate is involved -- neurons. So whether the selection process occurs in the same brain or in different ones, it is all memetic (except that inter-brain replication involves the communication problem identified by Sperber above).⁵

Indeed, a great divide separates the Mental Darwinists (usually motivated by evolutionary thinking) from the Intentionalists (usually psychologists or social scientists). Intentionalists don't see any way to avoid issues of meaning when describing human social activity, while the Mental Darwinists argue there is no need to engage in this intentionalist subjectivism to understand memetic processes. Any mental selectionist would prefer to naturalize psychology rather than make the many fine distinctions concerning the motivation behind information transmission (such as Conte's elaborate typology in this volume), regarding them as irrelevant to social dynamics.

Sperber agrees with Conte that it is absolutely crucial to distinguish between the commonality of beliefs, values and emotions that arise through transmission, and those that result from shared individual experiences (such as being in an earthquake) which do not involve any exchange of information between individuals. So the need for causal mechanisms that get information from point A to point B is clear. What remains unclear is whether this necessitates a turn toward intentionalism, or what Dennett (1971) has called the "intentional stance" (assigning beliefs and values as mental states to others). Perhaps such language is simply a necessary short-hand for discussing psychological processes in big-brained creatures, but should always be understood to be grounded in a Darwinian selection machine at the implementation level.

Social science

To the most pessimistic contributors, both social anthropologists (Kuper and Bloch), memetics is -- at best -- a promise at present, with no real results to show for itself. The question for these critics is whether memetics will ever contribute anything new to the explanation of society. For a variety of reasons, these anthropologists believe the answer to this question is no.

Ignorance of history

The primary reason for their cynicism is that they believe a quasi-epidemiological approach similar to memetics is already in widespread use in the social sciences, and indeed has a long (but

undistinguished) history in those disciplines. In their chapters, both social anthropologists go through an historical account of socio-cultural theory in anthropology to argue that memetics is old news -- and what is more, bad news. In particular, the idea that some cultures are more stable, or produce a higher quality of life because certain ideas spread better than others, has long been around. Thus, existing explanations for steadfast traditions and similarity in beliefs and values exist which do not invoke memes. However, such evolutionist approaches have been discarded, and superior theories have superceded them. Memeticists miss this "Big Picture" because they are largely unaware of the comprehensive literature which has accumulated in anthropology concerning cultural change, or the actual history of earlier views such as the cultural diffusionists of the early twentieth century (Bloch's target in particular). Being ignorant of the history of diffusionist thought in the social sciences, memeticists are simply condemned to repeat its mistakes.

What remains particularly unclear to these critics is the central claim of memetics: whether there is a novel replicator-based process underlying the population-level, epidemiological dynamic that is culture change. The primary problem of memetics, from this perspective, is whether there is a new entity on the horizon in whose interests things can be said to happen (the "meme's eye view"). This replicator would introduce a new kind of functionality which a social institution might serve: that of the memes. As such, it would represent a real and novel alternative to group-level functionalism, or the various flavors of structuralist thought current in the social sciences. Unfortunately, the central claim -- that a "meme's-eye-view" exists -- has not yet been proven.

These anthropologists also insist that there is a problem of circularity in memetic arguments. Memeticists only study things which seem likely to follow a memetic process, like fashions and fads. The perceived success of such empirical adventures leads memeticists to self-congratulation. But many aspects of culture aren't small, isolatable bits of information or practices that readily diffuse in observable time. Take the example of language, which permeates every aspect of culture. How does memetics expect to explain these more fundamental components of culture?

For some, even the word "meme" itself provokes problems. Its close parallel to "gene" may lead memetics astray, they argue, if in fact memes are not the same kind of thing. It also produces a "revulsion factor" among those who would otherwise be friendly to the Darwinian cause. Memetics is perceived by those outside the brotherhood as an arrogant usurper of territory, making extreme, unwarranted claims. This only serves to put memetics in the same basket with an earlier, related attempt at explaining human social life, sociobiology, which was widely seen as what Dennett (1995) calls "greedily reductionist." Sociobiology left no ground for social scientists to stand on, and all the interesting questions were subsumed under a single algorithm: the maximization of biological fitness. This is unpalatable to social scientists not just because of territoriality disputes, but because such reductionism is bound to failure. Thus, an undercurrent in the somewhat scornful reaction of even sympathetic social scientists to memetics is the perception that the social sciences will be "preempted" (Rosenberg 1981) by these evolutionary theories.

But in fact, this threat does not exist, as Plotkin (this volume) is anxious to point out. Can all social processes really be reduced to selection and transmission? The box of concepts available from Darwinism doesn't impress these social scientists. Memetics seems to employ a very small toolkit when so many theoretical alternatives are already available and there is so much complexity to explain. In fact, theory abounds in the social sciences. What is lacking is insight into real social processes. Explaining these seems a goal quite far removed from the concerns of most memeticists, who are laboring much further down the organizational hierarchy, worrying about replicators. An uphill battle against a wide variety of other approaches therefore lies ahead for memetics in the social realm.

In sum, memetics is seen as simply another case of those from outside the discipline, in this case largely biologists, "having a go at explaining culture," but without taking into account many of the complexities this project is widely recognized to entail. The meme critics are happy with the general notion that cultural change involves the diffusion of some vaguely characterized entity, but not with an explanation couched solely in terms of the selection, variation and inheritance of a particulate replicator.

Confusions about culture

The truly dismaying consequence of this critique is that -- as the social scientists themselves admit -- they do not have a viable alternative account of cultural change. What memeticists, in their general ignorance of social theory, also do not recognize is that the concept of culture -- the very thing which memetics intends to explain -- is itself sufficiently problematic that some social scientists advocate its abandonment. In their view, the notion simply covers too complex and varied a set of processes to be useful. (What exactly would replace the concept of culture, or what sub-concepts it should be divided into, is not obvious, however.) So, in some sense, the explanatory target memetics aims at -- culture -- is disappearing into thin air like the Cheshire cat.

At the same time, the anthropological enterprise is itself in serious trouble. So the question suggests itself: Is the major problem with the notion of memes themselves, or with the target it is meant to explain: culture? Those who take culture seriously, as the social scientists here do, find it hard to pin down their own conceptions of this central concept. One can only speak of an impossibly complex tangle of beliefs, behaviors and social institutions, as well as psychological predispositions and emotions, distributed through all the members of a society. Because all these things are linked, no possibility of reduction is admitted. As a consequence of this conceptual confusion, the project of explaining short-term cultural changes has in large part been abandoned by contemporary anthropologists.

But anthropologists admit that culture is distributed. If we can agree that much of cultural knowledge is socially learned, this implies that such knowledge necessarily diffuses through populations, from individual to individual. All sensory modalities require inputs in the form of temporal streams of information -- such as words forming sentences, and sentences paragraphs. At this basic level, individuals therefore must acquire information in bits (which need not be binary). So something like a unit of transmission must exist. If we can't speak of culture as a phenomenon that can be isolated, then perhaps we can still talk about the problem of how the ideational components of culture are learned through the social transmission of stimuli. The question then becomes how these units of transmission become translated or incorporated into the body of knowledge and practice that is culture. This is, in fact, Sperber's question -- and that of all the psychologists who assert that the psychology of transmission or communication is black-boxed by memetic evolutionists.

So regardless of the complexity of "culture" as a psychological construct in each person's mind, or as a set of practices and institutions, the informational components underlying the behavioral commonalities of culture (even in the standard anthropological view) must go through channels, migrating from mind to mind. And culture in this form -- if you like, exposed to the air as a stream of words, for instance -- can be studied in its own right. Indeed, the transmission process -- the foundation of how culture, in all its manifestations, is maintained -- is the rightful domain of memetics. How the bits of cultural knowledge get reassembled once they have reached a new, impressionable mind -- another fundamental process -- is the rightful domain of psychology. (But, as I have argued earlier, it is also one with important clues about how information acquired culturally is transformed before being sent out again into the social realm.)

The open question is whether these bits of information acquired through social transmission can themselves influence the likelihood they will be further transmitted. Do acquired units of exchanged knowledge have causal efficacy in human affairs independent of the wills of people themselves? In other words, are there memes?

Many arguments in the social sciences still center around the question of "agency," or the levels of causation. As Holy and Stulchik (1983:2) put it, the question concerning the level at which human behavior is caused is:

“basically about the autonomy of agency: if society, or structure, is an objective reality to whose demands people respond in specific ways, then it is an autonomous agency and individual people are its agents, and the only acceptable explanation is in terms of the functioning of the [social] system. If, on the other hand, society or structure emerges from, and is maintained or changed only by what people do, then individuals are autonomous agents and systems are the consequences of their actions and, in the last instance, explicable by them.”

This question -- individual or group -- has been at the center of the scientific status of social science since its beginnings -- with Durkheim, for example, falling on the top-down side of causative directionality, while methodological individualists, such as Rosenberg (1985), fall on the bottom-up side. Dawkins (1976) added a new, “lower” level of agency to biological theory by emphasizing that adaptations might reflect the interests of genes rather than individuals or groups. Similarly, Dawkins’ original meme suggestion indicated that a new, lower level of agency might also be relevant to the explanation of social facts. The “meme’s eye view” shifts the location of cultural agency below the standard “floors” of individuals or groups to the “basement” level of information itself. However, such a hoary controversy as that concerning the location of agency is unlikely to be settled here. And even if replication is found to underpin some cultural knowledge acquisition, it is unlikely to be the whole picture, as Sperber (this volume) argues. So it is improbable that memetics will ever provide a full account of cultural change; some aspects of cultural continuity will be due to the push and pull between genes and environment.

There has been an implicit dualism of the agency debate, with the available alternatives have typically been presented as an “either/or” choice. That is, either individuals are assumed to be fully independent agents, or individuals’ cultural repertoires are thought to be fully determined by the society in which they live. A similar restriction has also infected the debate about memes. But in fact it seems likely that individual learning direct from the natural (exogenous) environment can co-occur with social learning, both from other members of the society, as well as from capitalized resources such as books. I think Laland and Odling-Smee’s concept of environmental inheritance through niche construction goes a long way toward handling the additional complexity of culture as outlined by Kuper and Bloch. The “built environment” (including technologies for information storage such as books and computers) which certainly constrains human action, is after all a consequence of the activities of previous generations. Having three forms of inheritance (genes, memes and artefacts) is a means by which a sophisticated theory of mutual constraint relations between individual, societal and cultural replicator levels can be constructed within an explicitly evolutionary framework.

Progress in memetics?

In the general struggle to understand culture, there is a clear trend for increasing divergence between groups, with decreasing mutual intelligibility. One line is becoming centered around cultural studies, while the other seeks refuge in science. Memes are perhaps more and more likely to be the rallying cry for Darwinists of all stripes when discussing culture, while simultaneously being an object of derision among those inspired by the humanities. Memetics may thus play its small part in the increasing division among researchers. Perhaps the debate is not really about memes at all, but rather more a matter of temperament than anything else. At bottom, whether you “like” memes may be simply due to whether you are a splitter or a lumper, a believer in analysis or interpretation.

Although it is obvious that despite the shared belief among those collected here that some kind of evolutionary approach to culture is necessary, significant barriers to communication remain between those from different disciplines. This perhaps derives from the varying histories these disciplines have had with evolutionary approaches. Biologists are predisposed to look at issues of transmission because inheritance is central to their subject, while those trained in the social sciences have been more interested in structure and function -- which have traditionally been answered without attention to dynamics, much less the more specific question of transmission. Nevertheless, social anthropology has a long history of evolutionary thought, broadly speaking, which has generally not proven successful. Indeed, a common refrain among social anthropologists seems to be “been there, done that.” It will be

difficult for believers in memes to convince these historically mindful and hence reticent social scientists that this time around things might be different. Similarly, it has proven difficult for the anthropologists to explain exactly what went wrong with previous incarnations of cultural evolutionism, or specifically how the memetic perspective is likely to go wrong itself, even if given a clear run at explaining culture.

But other factors besides academic background also seem to dictate use of the word “meme” in scientific circles. For example, the teams of Boyd-Richerson and Laland-Odling-Smee both use the same formalism for investigating cultural evolution. But one team rejects while the other accepts the idea of particulate, transmissible units of information as necessary components of the explanation of culture. While Boyd and Richerson may be more enamored of the theoretical possibility of non-particulate inheritance, Laland and Odling-Smee appear to be more impressed by the need for replication to effect transmission. Other rejections of memes are probably idiosyncratic or, perhaps, reflect the continuing confusion surrounding the word “meme.” Given this multi-layered resistance to memetics, it may be wiser to follow the progress of evolutionary cultural studies more generally, rather than the meme idea *per se*, for a true indication of who is winning the battle to explain culture.

Applying memetics

The question of whether memetics has an empirical future remains open. Among partisans and detractors alike, a major disappointment with the current status of the field is the lack of studies in what might be called “applied memetics.” Hull (this volume) argues that we should all just go out and “do it.” However, it is not clear that such an approach will be successful if I am correct about the need to identify the responsible mechanisms underlying cultural inheritance. Instead, I would suggest memetics must first establish how cultural traits maintain themselves in similar forms through generations. Perhaps many mechanisms will be involved, as it is possible there will be as many mechanisms as there are media for social learning.

So we need to develop specific methodologies for conducting memetic studies. There should also be more discussion of existing empirical studies that weren’t undertaken under the banner of memetics but which could be interpreted as falling within the general purview of this incipient discipline.

It may be that it won’t be possible to conduct empirical research in this area for the simple reason that the process being investigated is too complex. From my own experience (Aunger 2000), I would suggest that the prospects for fruitful empirical studies in memetics are daunting. Despite dogged concentration on a highly restricted question (transmission of a limited set of beliefs in a “simple” oral society), and the application of various multivariate statistical techniques, I have been unable to provide a quantitative estimate of the relative strengths of intra- and inter-generational transmission. On the other hand, a rather more limited transmission science may be possible -- and valuable. For example, the exact magnitude of selection coefficients are often unknown in biological studies, but also without much interest: what we really want to know is whether selection is directional rather than neutral, and to identify the selecting agent. The answers to these kinds of questions can get us a long way toward an understanding of the evolution of the system under study and may be possible for a future memetics.

At any rate, as even David Hull (this volume) acknowledges, given the extensive theoretical work already accomplished and the high level of current interest in the subject, something substantial can rightfully be expected of memetics in the relatively near term -- either by way of correct, novel predictions derived from the meme hypothesis, or proof that cultural entities with the characteristics of replicators exist. This is because the ultimate test -- which would preempt theoretical objections -- is whether memetics can produce novel empirical work or insightful interpretations of previous results. It has not yet done so, but must do so in the near future. Otherwise, it is likely that memetics will be perceived to be a misguided enterprise. The clock is ticking.

Acknowledgements

Thanks to Gillian Bentley, Rosaria Conte, Liane Gabora, David Hull, India Morrison and Henry Plotkin for feedback on an earlier version of this chapter.

Footnotes

1 This is not the conclusion Laland and Odling-Smee necessarily intend us to draw when they assert niche construction can be important, but their framework can switch the burden of explaining technology from cultural to ecological inheritance.

2 Some (e.g., Blackmore 1999) see the development of human language as an early example of gene-meme cooperation: “digitized” signalling achieves increased fidelity for memes during transmission, while grammatical structure allows increased sophistication in message-passing for human social coordination. But language is subject to the problem associated with any form of communication: the need to transform brain-language to a public code and back again. See the section on memetic phenotypes for a discussion of this difficulty.

3 Actually, there are various models of replication within brains -- see Delius 1991; Calvin 1997; Aunger 1999 -- but none that work at the neuroscientific level between brains, which is what I’m talking about here.

4 Given this basic level of uncertainty about the nature of memes, it seems to me premature to start making distinctions among memes as a number of authors have done. Plotkin’s distinction (in this volume) between “surface-” and “deep-level” memes is similar to the standard one in psychology between procedural and descriptive knowledge, or, roughly speaking, between a knowledge of things and how to do things with things. Scott Atran (1998) recently distinguished between “core” and “developing” memes. Core memes are acquired through informationally-encapsulated modules designed by natural selection; developing memes fall into the cracks between modules, and therefore having to be processed by some amalgam of processing units. Core memes therefore last longer, are more reliably acquired, and generally have the desirable features of good replicators, in Atran’s opinion. Such distinctions depend not only on a good knowledge of how encapsulated information processing algorithms are, but also about how memes might interact with this mental architecture. This makes such propositions doubly “courageous.” I think we first need to establish that there are memes (as mental entities) before we start dividing them into species (Aunger 1998).

5 An unfortunate consequence of a selectionist psychology is the appearance it gives of there being no room for human agency in decision-making; that all human psychology is merely a random selection process among alternative behavioral choices. But of course the abandonment of intentionality and free will would be hailed as a victory for memetics by hard-core Mental Darwinists.

References

- Atran, S. (1998). Folk biology and the anthropology of science: Cognitive universals and cultural particulars. *Behavioral and Brain Sciences* 21:547-569.
- Aunger, R. (2000). The life history of culture learning in a face-to-face society. *Ethos*.
- Aunger, R. (1999). Culture vultures. *The Sciences* 39:36-42.
- Aunger, R. (1998). The 'core meme' meme [Comment on 'Folk biology and the anthropology of science: Cognitive universals and cultural particulars' by Scott Atran]. *Behavioral and Brain Sciences* 21:569-570.
- Benzon, W. (1996). Culture as an evolutionary arena. *Journal of Social and Evolutionary Systems* 19:321-362. [<http://www.newsavanna.com/wlb/CE/Arena/Arena00.shtml>].
- Blackmore, S. (1999). *The meme machine*. Oxford: Oxford University Press.
- Calvin, W. H. (1996). *The cerebral code: Thinking a thought in the mosaics of the mind*. Cambridge, MA: MIT Press.
- Changeux, J-P. (1997 [1985]). *Neuronal man : The biology of mind*. Princeton: Princeton University Press.
- Dawkins, R. (1982). *The extended phenotype*. Oxford: Oxford University Press.
- Dawkins, R. (1976). *The selfish gene*. Oxford: Oxford University Press.
- Delius, J. (1991). "The nature of culture." in *The Tinbergen legacy*, (eds. M. S. Dawkins, T. S. Halliday and R. Dawkins), pp. 75-99. London: Chapman and Hall.
- Dennett, D. (1995). *Darwin's dangerous idea*. Hammondsworth: Penguin.
- Dennett, D. (1971). Intentional systems. *Journal of Philosophy* 68:87-106.
- Gabora, L. (1997). The origin and evolution of culture and creativity. *Journal of Memetics - Evolutionary Models of Information Transmission* 1. [http://www.cpm.mmu.ac.uk/jom-emit/vol1/gabora_1.html].
- Gatherer, D. G. (1998). Why the thought contagion metaphor is retarding the progress of memetics. *Journal of Memetics - Evolutionary Models of Information Transmission* 2. [http://www.cpm.mmu.ac.uk/jom-emit/1998/vol2/gatherer_d.html].
- Hallpike, C. R. (1979). *The foundations of primitive thought*. Oxford: Oxford University Press.
- Holy, L., and Stuchlik, M. (1983). *Actions, norms and representations: Foundations of anthropological inquiry*. Cambridge: Cambridge University Press.
- Reader, S. M., and Laland, K. N. (1999). Do animals have memes? *Journal of Memetics - Evolutionary Models of Information Transmission* 3. [http://www.cpm.mmu.ac.uk/jom-emit/1999/vol3/reader_sm&laland_kn.html].
- Rosenberg, A. (1985). *Philosophy of social science*. Boulder, CO: Westview Press.
- Rosenberg, A. (1981). *Sociobiology and the preemption of social science*. Baltimore: Johns Hopkins University Press.
- Tooby, J., and Cosmides, L. (1992). The psychological foundations of culture. in *The adapted mind*, (eds. J. H. Barkow, L. Cosmides and J. Tooby), pp. 19-136. Oxford: Oxford University Press.