

Three Roads to Cultural Recurrence

Robert Aunger

(in press, *The Epidemiology of Ideas*, ed. by Dan Sperber, Chicago: Open Court)

Social scientists have long remarked that there is consistency in what people believe and value over time, especially within definable groups. Anthropologists call this body of information “culture.” There are (at least) three causal mechanisms that can explain the recurrence of cultural traits. Recurrence can occur through 1) strong individual learning biases; 2) population-level normalizing effects on what is adopted; and 3) replicator-based inheritance. Each of these mechanisms is favored by a particular brand of evolutionary theorizing about human society. Evolutionary psychologists (EPs) advocate the first option, which emphasizes the ability of universal structures in the evolved mind to come up with the same responses to environmental conditions time and again. What explains cultural consistency over time, then, is evolved psychological decision-making processes in the face of common environmental challenges (Tooby and Cosmides 1992). A group I call “cultural selectionists” (CSs) prefer the second option, which notes that even poor social learning abilities can still produce consistently shared features at the level of the group if there are widely shared psychological preferences for traits or the types of individuals from whom to acquire culture (Boyd and Richerson 2000; Henrich and Boyd in press; Gil-White 2001). The third option, based on replication of the same information from generation to generation, is the memetic position (Dawkins 1976; Blackmore 1999; Aunger in press a). In this scenario, the cultural features that keep popping up are the phenotypic expressions of memes, or cultural replicators, disseminating through the population via social communication. This variety in the possible explanations for cultural evolution is not generally recognized, nor do advocates of one position generally acknowledge the validity of others. But I will argue in this paper that all three of these possibilities are viable in our present state of ignorance about the means through which cultural traits reappear each generation; any one of them may account for a particular aspect of cultural inheritance.

First road

Getting ideas into other people’s heads is a problem. People are like islands, isolated in space. Neither bits of brain (neurons) nor infobits of “brain code” (the action potentials neurons use to communicate to one another) survive in the environment outside the brain. All human communication is thus mediated in the sense that a “public” code must be created that two brains can share (Sperber and Wilson 1986; Sperber 1996, 2000). Social communication is thus always achieved through the use of signals, acting as intermediaries. These signals can be generated by bodily movements like hand-waving or the movements of lips and palate accompanied by the forceful expression of air (speech “behavior”), and so on. In modern Western societies, signaling is also increasingly made more complex by the intervention of artifacts: the signals we consume have been modified through interaction with books or computers. People make the artifact; the artifact in turn modifies “ambient” signals (such as reflected sunlight) or generates new ones (such as electricity through wires), which become what we interact with.

From an evolutionary viewpoint, the problem is that information is constantly being reconstructed in the course of communication – by receiver’s minds, and possibly by interaction with artifacts prior to reception; what we eventually acquire as a message has been

through a suite of transformations. As widely acknowledged in the literature on human communication, signals are “impoverished” in informational terms – not everything that the sender wants to communicate is present in the signal itself; there is a difference between message and meaning. In effect, successful communication depends on the receiver’s ability to correctly *infer* meaning from message. Why are messages impoverished? Analogy can be made to the “translator’s problem” of recreating a poem written in another language: meaning, and nuances of implication, are almost inevitably “lost” in the process. Similarly, we can assume that the mere act of translating from brain code to public code results in some loss of information. Second, we simply aren’t aware of everything we wish to communicate; some of our message remains subconscious. Part of this subconscious information no doubt gets communicated through non-verbal means – bodily attitudes, tones of voice – but presumably there are intrinsic limits to the kinds of information that any given channel can bear.

But it is commonly observed that similar traits tend to pop up generation after generation, particularly within groups that communicate with one another, and which often identify themselves as sharing a “culture.” If we are to explain the recurrence of cultural features in the face of this sloppy transmission of information, we need to postulate another factor that keeps culture “on track.” For prominent evolutionary psychologists such as Steven Pinker (1997) and Geoffrey Miller (1999), this mechanism is tightly channeled individual learning. Shared, evolved mental routines for processing environmental stimuli assure regularized responses to recurrent events. EPs thus argue against the commonly-held conception that culture is largely transmitted through social learning from individual to individual. Instead, it is the universal ability to derive the proper response from features of the environment that ensures the similarity of beliefs and behaviors within the groups that share those environments. If social transmission is unreliable, and cannot guarantee that individuals have the information they need when a dangerous occasion or important opportunity arises, then evolution would have opted for a more secure option. It would have placed the relevant information inside each of our heads from the beginning, through genetically inherited content and structures for judiciously ascertaining just what kind of situation any given state of affairs represents. “Culture” is thus “evoked” in people from a repertoire of stored memories by their similar surroundings and common mental algorithms. These mechanisms evolved over time for their ability to produce optimized responses in evolutionarily important situations.

In particular, EPs have little use for cultural dynamics. Steven Pinker is scathing on the idea that cultural evolution introduces a new kind of evolutionary force into human life. He says (1997:210) “Nothing in culture makes sense except in the light of psychology. Evolution created psychology, and that is how it explains culture.” For Geoffrey Miller (1999), culture is not a system of inheritance in which useful knowledge and group-benefiting traditions are passed through the generations; rather cultural behaviors constitute courtship displays in which individuals try to attract and retain sexual partners. In neither case is there an admission that culture can lead to emergent phenomena such as social institutions that can feed back on human psychology.¹

Second road

However, the EPs are mistaken in their belief that because transmission is inherently error-ridden, it cannot underlie cultural evolution. The EPs’ mistake is to think that only a gene-

¹ Tooby and Cosmides (1992:121-2), however, are more eclectic; they admit that social institutions and maladaptive traditions can emerge at the group-level through epidemiological information dynamics.

style approach to transmission will guarantee accumulated advances, as well as fidelity to valuable past traditions – cultural recurrence, in effect. All of the virtues we find so admirable in genetic evolution derive from the accurate, unbiased replication of genetic material. Mutations rarely occur as DNA molecules are copied, thanks to multiple safe-guards on the process, and individuals are just as likely to get their mother's as their father's gene for any given copy, so there is no intrinsic sorting of genes as they are inherited (Henrich and Boyd in press). If cultural evolution is to have these qualities, the EP thinking goes, then it will have to be based on the same mechanism as the genetic case: unbiased, error-free replication.

However, Robert Boyd and his colleagues (Henrich and Boyd in press; Boyd and Richerson 2000; Gil-White 2001) have recently shown that cultural evolution need not exhibit unbiased, error-free replication for that process to exhibit cumulative adaptive outcomes. There is in fact more than one way to skin a cat, more than one road to Rome. Adaptive culture can evolve even if transmission is sloppy. The key to accumulating adaptations is the possibility that cultural traits are *non-randomly sorted*: individuals can use their minds to preferentially adopt particular variants from among the traits circulating in the population. This gives a kind of directionality to cultural evolution missing from its genetic cousin, with its Mendelian rules of inheritance. Any such predilection to adopt particular variants we can call a psychological "bias." These biases may themselves be just the sort of thing that EPs talk about: universally shared, evolved modules for processing information in certain substantive domains in specific kinds of ways. The bias can be "direct" in the sense of depending on the content of the trait, or "indirect" in the sense of reflecting a preference for other features of the trait, such as who is harboring it. If we assume that there is a bias toward imitating individuals with high social prestige, for example, or even just a predilection toward adopting the most common trait in the population, cultural "inertia" can be exhibited by the population – there is a strong tendency for the dominant trait to remain dominant, and to continue to characterize the population, even in the face of a very high rate of error in transmission (that is, a relatively low likelihood that a given individual will infer the correct trait value from his or her learning experiences). In effect, many instances of the same bias lead to a cumulative reduction in the measured error of transmission at the population level; error is absorbed by being consistently biased in one direction because everyone shares the same cognitive bias. This gives a directionality to changes in the frequency of traits that wouldn't be apparent from looking only at the apparently random errors that individuals in that population are making.

So cultural evolution can be both error-ridden and biased and still accumulate adaptations, cultural selectionists assert. Highly biased transmission can effectively compensate for a high rate of mutation in cultural inheritance. Even if the trait is characterized as a quantitative distribution rather than as having discrete values, this tendency holds true: the population can maintain a cultural tradition despite strong proclivities among individuals to make mistakes about what they should be learning from others in the group. As long as the psychological bias is sufficiently strong relative to selective forces, representing cultural variation as a continuous spectrum of values has no effect on the outcome: the weak selective component determines the equilibrium values of the population (Henrich and Boyd in press). As Henrich and Boyd note, population genetic models were providing insights into the nature of biological evolution before anyone knew anything about DNA discreteness or replication. Further, genetic models of quantitative characters, like height, still ignore these "units." So adaptation occurs even if you characterize cultural variants along a gradient. Further, mathematically speaking, the outcome is the same as if the process had been that of replication in discrete traits. That is, the population winds up with the same cultural configuration as would occur had memes been replicating through communication between individuals all along.

Henrich and Boyd (in press) also find that cumulative adaptive evolution is facilitated when the social group is larger – as might be expected when they are basically aggregating over

errors to achieve their desired result. Things also work out better when individuals tend to make independent “mistakes” in their learning, so that there is no consistent pressure working in a direction other than that of selection. No matter how poor individuals are at imitating one another, cumulative evolution can still be generated for some combination of error-proneness and selective pressure. Cultural selectionists conclude that we must consider both the social environment and individual cognition in our thinking about cultural change because both influence the conditions for cultural adaptation.

We can thus conclude that it is an error to equate evolution with replication – even though one of the premier contemporary evolutionary theorists, Richard Dawkins (1976, 1982, 1999), continues to assert that evolution depends on replicators. As Lewontin (1970) pointed out, a process need only exhibit inheritance, selection and variation in order to qualify as evolutionary. You therefore cannot argue, as the EPs do, that cumulative cultural evolution is not possible because cultural transmission leads to different representations in each individual’s mind. Replication is but one mechanism of inheritance, so evolutionary theory encompasses a larger universe of possibilities than the example of DNA suggests. Culture can evolve through other means than genes do.

Dan Sperber and his colleagues have suggested a specific mechanism that produces cultural recurrence without replication, and which is therefore consistent with the CS position. They begin with the need to reconstruct meaning from messages, as pointed out at the beginning of this paper. In particular, mental representations in person A give rise to behaviors (which Sperber calls “public representations”) that are observed by person B, who must then infer A’s mental representation from that public version of the information. Sperber (1996 and elsewhere) has encapsulated this perspective, with its sequence of translations from private to public representations, in the phrase “the epidemiology of ideas.” Further, because individuals differ and public representations provide incomplete information, this inferential process can be highly inaccurate (Sperber 1996:103-118; Atran 2001). However, in some cases, psychological attractors work to ensure that sender and receiver meanings are sufficiently similar to call them replicas of one another (Sperber 1996:104). In effect, variations introduced into a message, as inferencing mechanisms to ascertain the meaning of a communication, must either stay close to “home” or cancel each other out. Cases of attraction, where similarity of mental representations *does* result from communication, may reflect the operation of evolved modules, as postulated by EPs. In other cases, these modules can themselves be the products of the cultural exchange of information. But replication is just the limiting case in which no transformation of meaning takes place during communication. In general, we need a social science of culture that takes the special features of cultural recurrence into account – particularly the different possible outcomes of signaling, which range from replication to great disparity between mental representations. What needs explaining is how mental representations can sometimes recur stably over generations in a considerable proportion of a group’s population, to constitute their “culture,” given the tendency for inferencing modules to go astray.

Sperber thus agrees with the cultural selectionists that replication is not the only road to cumulative adaptation. Like CSs, he also fully accepts the modularity hypothesis of EP. His difference from these other positions arises in his commitment to the hypothesis that psychological attractors help cultural lineages persist in the face of messy social transmission.

Third road

There is a third road which none of the parties discussed thus far care to travel, and that is the memetic road. Although perhaps the least traveled, I believe it still remains a viable way of

winding up in Rome. The cultural selectionists have saved the idea of cultural transmission from the onslaught of the evolutionary psychologists, but at the cost of losing the possibility of replication as a consequence of that transmission. My argument is that there is a way to save both the baby and the bathwater, a way to mix heavy-duty mental reconstruction with the social replication of information. This may seem an impossible combination, given the arguments I have just rehearsed, but let me attempt to justify this position in the following.

Sperber suggests that communication is about the sharing of thoughts, about the recreation of a particular message or meaning in the minds of others (Sperber and Wilson 1986:1). But to an evolutionist, this is only a proximate goal for communicators; the ultimate goal of communication is to manipulate the mind and hence behavior of others through the manufacture or display of signals or signal-generating artifacts (Dawkins and Krebs 1979). For example, the sender may need help to achieve some goal that can only be achieved through cooperation, like moving a piano, or raiding an enemy's camp. But inducing cooperation may not depend on a duplicate of the sender's thought being produced in the receiver's mind. Instead, deception or irony may be involved, in which case the speaker may not want his or her true intention to be inferred, and hence duplicated in the receiver's mind. People depend on their mutual abilities to mind-read to help achieve this goal. Mutual mind-games are a major part of sharing culture. So replication is not a prerequisite to successful communication, just as the EPs and CSs claim.

But at the same time, EPs posit that there are heavy-duty regularizing structures in human brains that ensure that inferences take a certain form, given a stimulus within some domain defined by content (such as ethnobiology, naïve physics, or theory of mind). These are evolved mental structures that presumably all people share. (As I have just argued, CSs suggest that minds contain evolved biases, if not encapsulated modules, as well.) This would seem to provide everyone with a kind of error-correction mechanism that guarantees that, despite the poverty of the stimulus, the "correct" inference is nevertheless routinely drawn from a signal or message. Thus, the representation that winds up being constructed in the mind of the receiver is very much the same as that which produced the signal in the first place in the mind of the sender, because these coding and decoding mechanisms are shared. (I ignore the additional step of inferring sender intent because it may be irrelevant to the message being successfully replicated.) If inferencing is regularized by evolved predispositions to treat certain kinds of input in a similar manner, then the conclusions reached – even based upon the relatively poor information of a signal – may be stably reproduced on multiple occasions, in which the same stimulus results in the same conclusion by different brains housing the same inferencing machinery. A set of similar mental representations can then be produced by a chain of events in which the same signal is reproduced over and over by a sequence of individuals.

What exactly must hold if replication is to be true of cultural learning? In Auger (in press a), I argue (based on literature in the philosophy of biology and Sperber 2000) that replication is a special relationship between a source and a copy such that four conditions hold:

- causation (the source must play an active 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).

Based on this model of replication, I argue (in Aunger in press a) that certain specialized kinds of interneuronal communication could lead to the replication of information within a brain.

As Sperber emphasizes, signals don't contain thoughts; thoughts stay inside our heads; they are private representations of things in the world or things that are purely imaginary. We have to think about the nature of signals and what they can accomplish if we are to understand replication in the context of communication. This "new view" must find a role for signaling that is consistent with significant local reconstruction of information in the recipient's brain. Hints come from Atran (2001:353), who calls signals "elicitors that draw out inferences and information from the mind;" similarly (and independently), I call them "instigators" (Aunger in press a). The role of signals in this view is not to bring along the information that will be incorporated into the meme copy, but rather to set in motion the kind of local transformation that will result in the copy being produced.

This is a different way of thinking about replication, but one which is consistent with what we know about how human communication works, particularly the "poverty of the stimulus" tenet. The social learning process no longer need be based on imitation. As Atran (2001) has pointed out, imitation is not only the cause of replication, but also represents the information to be replicated – a kind of self-contained form of information transfer. The message contains sufficient instructions for copying itself – a form of self-referentiality. This "complete package" assumption, associated with the traditional approach to memetics, and with its obvious analogue in the form of genes, appears to contradict the poverty of the stimulus assumption. In Aunger (in press a), I argue this poverty could be complemented by a "richness in the response," the ability of mental machinery to reconstruct just what the stimulus "says," regardless of how little information is actually present in the stimulus.

What has not been generally noted is that, at least theoretically, signals and meanings can be informationally independent. That is, the information present in a signal may bear precious little relation in pattern or content to what is present in the mind that generated the signal in the first place. In effect, the copy of the message need not incorporate information from the signal, which only has the job of starting off the message inferencing process. This view of a signal's role as instigator has the virtue of making memetics Darwinian rather than Lamarckian. If we can get the replication of information largely through regularized local reconstruction, then what happens to signals becomes irrelevant. As long as signals can instigate the reconstruction process, that is enough to ensure replication of the message, thanks to the internal guidance and error-correction systems of the mental module itself. Any mutations introduced into signals through their travels do not get incorporated into the "genotypic" line of mental replicators themselves. This is the cultural equivalent of the central dogma of molecular biology, that phenotypic information does not get inherited (a proposition which Lamarckianism denies).

This view also suggests that there is no reason to suppose that significant error will accumulate in a cultural lineage. The idea that information replicates successfully, even in the face of significant reconstruction, is thus preserved – and in fact promoted by the notion of universal mental modularity. Those potential elements of culture that are most reliably reproduced by people's evolved, modularized minds are those that recur and persist, just as evolutionary psychologists argue. It just happens that their proximal cause is social signals rather than learning from the environment.

The only question is whether, in the case of "heavy-duty" reconstruction, there is informational inheritance -- is the copy similar to the source because information deriving from that source has been incorporated in the copy? If individuals are largely recreating cultural knowledge for themselves, based on the flimsy evidence of what others signal, this

seems to violate the inheritance requirement of an evolutionary system that information be passed from individual to individual in a lineage (as parents “communicate” their genes to their offspring). Sperber (2000) argues that if the reconstruction can be said to be due to genetically inherited modular inferencing mechanisms, then cultural transmission is not the primary cause of cultural similarity between people. Rather, each person individually relearns much of that knowledge thanks to naturally evolved modules for inferring things. The cultural link may appear to be broken by this “on-board” constructivism. But signal-based instigation does satisfy the inheritance condition: the message the signal brings is what causes the desired result in the recipient brain – replication of information similar to that in the source brain. No one has ever argued that replication was achieved without assistance. Even genetic replication requires the participation of many agents besides DNA (like messenger RNA and ribosomes). So the fact that cultural knowledge is inferred, based on the content of signals, does not exclude it from the category of replication events. And the fact that inference is commonly regularizing or normalizing suggests that it is quite often the case that the duplication of cultural knowledge occurs through this transmission process. This argument would seem to invalidate one of the primary criticisms of the EPs, that variation occurs each time information is passed from person to person, leaving memetics little to explain.

The view of memes I have suggested is also consistent with another claim: that those ideas which do disseminate well are those which make good use of evolved domain-specific psychological machinery (Sperber 1996; Atran 2001). The memes that persist are those that “fit” with the evolved psychology; ideas which try to cross or slip between boundaries of evolved inferential categories will not be good replicators. Memetics can thus take on board much of what EPs have to say about the existence of genetic structuring of the brain without tossing out the central notion of replication. And since replication is one possible mechanism producing cultural inheritance, memetics can also be considered a specific hypothesis about how the cultural transmission of information acknowledged by CSs is achieved.

Conclusion

Evolutionary psychologists have argued that culture does not evolve through the social communication of information; instead, it is largely recreated anew by each individual, thanks to shared, evolved modules for inferring meaning from messages. In response to this claim, Boyd and colleagues have emphasized that evolutionary processes, including cultural ones, can retain the ability to accumulate useful variation and complexity through even highly inaccurate social transmission. In the CS view, transmission is an important component of inheritance; it just works differently than in the case of genes. In particular, it need not depend on replication as a mechanism. The cultural selectionists have proven their point that there is a real need to understand how individual predilections convert into population trends over time, and that we need formal models to assist our poor intuition in this regard. We can’t easily guess what the implications of a pattern of individual interactions will be in our heads. Counterintuitive social phenomena can derive from intuitive propositions about our native psychology. For example, we shouldn’t assume that the only process that can give rise to accurate reproduction at the level of the population is accurate replication at the level of individuals. Cumulative cultural trends at the population level can be achieved without replication at the individual level.

At the same time, Sperber’s work has had the salutary effect of keeping everyone honest, of recognizing peoples’ need to make the public/private coding switch – a fact too easily forgotten in the abstractions of the cultural selectionists’ mathematical modeling, where genotypes and phenotypes can be switched through a purely interpretive move, without a jot being changed in the formulas. Also, the emphasis on mental reconstruction has required

some rethinking about how cultural inheritance might work. Any viable theory of cultural evolution now has to be able to jump through the hoops that Sperber has put in place. This restriction of possibilities is just what Boyd and Richerson long for, to limit the Design Space they have to explore with their formalisms (Boyd, personal communication).

While we all agree that replication isn't a *necessary* feature of cultural evolution, it still may be characteristic of communication in some contexts. Here and elsewhere (Aunger in press a), I have argued that, contrary to both the EP and CS positions, you can have replication in the face of significant mental reconstruction. The theoretical work by Boyd and colleagues shows that memetic inheritance can be mimicked by other kinds of processes (such as universally biased transmission). Still, the existence of memes, if proven, would have significant repercussions: cumulative adaptation is much more likely across a broad range of circumstances if there are replicators at work in the inheritance of cultural traits.

The replication question is important because there are limitations to the conclusions reached by CS modeling. First, the CS models depend on everyone having the same psychological biases, and over a long period of time. Without a strong bias toward particular values, the population would soon become awash in random variation, and only very strong selection pressure would be able to maintain any directionality or consistency in cultural values over time; it is the shared social bias for certain learning models or trait values that compensates for individuals' proneness to make psychological errors. In the absence of such a bias, we get a very different scenario: near-zero cultural heritability.

Further, such biases aren't necessarily going to characterize many real-world instances. Certainly there is no general tendency to imitate prestigious individuals or favor particular cultural values in the case of one extensively studied cultural system, that of food taboos in the Ituri Forest of the Democratic Republic of Congo (Aunger in press b, 2000). Instead, individuals acquire their ideas about the edibility of particular foods from specific individuals designated for them by other normative rules about culture learning: children should learn food taboos from their father. This norm tends to produce significant variation between the cultural lineages defined by individual-to-individual transmission patterns. No global psychological bias causes the population to adopt consistent beliefs over time; there is, in fact, considerable intra-cultural variation, with each individual exhibiting a unique suite of taboos. If this is an indication of a more general trend for real-world cultural systems to be culturally rather than psychologically determined – and the number of things you are supposed to learn from your parents would suggest that food taboos in the Ituri are not the only example of such a system – then the selectionist models are more limited in their implications than CSs would like.

Second, as Henrich and Boyd (in press) admit, it is probably the case that the very fine control over replication in DNA allows a given selection pressure to be more effective over the long term, permitting the accumulation of very precisely adapted traits over time. On the other hand, the constant pressure of drift through errors introduced into what is socially learned washes away some of this precision in cultural evolution. Cumulative adaptation becomes doubly problematic if there is a tendency for psychological biases to shift over time; if biases are not consistent from generation to generation, then no cumulative directional psychological selection of traits will take place. This implies that such biases must not only be universally shared, but slow to change (as might be true of a gene-influenced bias). If biases are themselves culturally acquired, then they can become fad-like, and cumulative cultural adaptation is lost again.

Henrich and Boyd (in press) are also enthusiastic about quantitative models of cultural inheritance; they feel this possibility liberates us from the need to think about replication at all. In their view, we can think of culture as a continuous spectrum of beliefs or values that

blends together in our minds, much like height in biology can increase in seemingly infinitesimal increments. This is a perspective which would appeal to some neuroscientists, who think of the brain as existing in a kind of complex global state at any moment in time which it is futile to tease apart since each part of its big network of neurons can influence every other (e.g., Freeman 1999, McCrone 1999). But at least in the biological world, the quantitative inheritance of traits like height is still underlain by qualitatively different traits: particulate genes. It is just that multiple gene loci are involved in the production of the relevant phenotypic spread. The effects of different genes are blended together to produce an organism's height, but the units of inheritance producing them are not blended together at all: each gene locus simply adds value to the height achieved by an organism. I think the same could be true of mental traits: the global state of the neuronal network will not be the relevant parameter for describing inheritance of a cultural trait; instead, particular segments of the network, working together, will be responsible for the information associated with that trait. Memes will be shown to be partible physical things.

Still, this is a conjecture. Whether social replication actually occurs remains to be discovered through empirical work – primarily in neuroscience. However, it is worth noting that we don't currently have an example of an evolutionary process that is not founded on a replicator – genes, prions, and computer viruses are all replicators, and together account for nearly all the unique evolutionary processes we see around us (Aunger in press a). (The “epigenetic” inheritance of biological information through cell types does occur, but this tends to account for only a small fraction of organic traits.) Maybe culture is going to be the first case of a major evolutionary process with a different kind of etiological mechanism; we just don't know at present. And even if cultural replication occurs through the kind of mechanism I hypothesize, it may only account for a relatively small proportion of shared knowledge, the rest being produced through independent learning rather than as a result of stimuli coming in from other members of the social group.²

Everyone acknowledges at least one contribution from the EP camp: the emphasis on shared learning algorithms or modules. Boyd and Sperber build on this foundation by hypothesizing that one of these mental modules may be evolved to process social signals, enabling the benefits of social learning to be widely shared. Sperber and the memeticists then add onto this foundation specific hypotheses about how social learning works. It therefore seems we have a much larger universe of possibilities for mechanisms of cultural evolution than has been generally recognized.

The question that remains is which view is correct. Possibly, each is a viable explanation for some kinds of cultural traits. Culture might be a heterogeneous body of knowledge acquired through a variety of routes. We just don't know at present. Three routes to cultural recurrence, each with unique features and therefore different evolutionary prospects, continue to be active prospects. The quest to find out which view is most appropriate is on, and of great interest, because which kind of inheritance mechanism turns out to characterize particular cases will tell us a lot about just what kind of evolution that trait will experience.

² In similar fashion, prions only account for a small proportion of diseases, but the peculiar pathologies they do cause – transmissible spongiform encephalopathies like “mad cow” disease – nevertheless represent an interesting species of disease which requires its own kinds of analysis and treatment. The same would be true of memetic culture: significant interest would attend the discovery of memes, and projects to uncover the replicative abilities of cultural traits and the unique kinds of dynamics they introduce would naturally develop.

ACKNOWLEDGEMENTS

A previous version of this paper was profitably read by Robert Boyd and Dan Sperber.

REFERENCES

- Atran, Scott (2001) "The trouble with memes: Inference versus imitation in cultural evolution." *Human Nature* 12:351-81.
- Aunger, Robert (in press a) *The Electric Meme: A New Theory of How We Think and Communicate*. New York: The Free Press.
- Aunger, Robert (in press b) "Exposure versus susceptibility in the epidemiology of everyday beliefs." *Journal of Cognition and Culture*.
- Aunger, Robert (2000) "The life history of culture learning in a face-to-face society." *Ethos* 28(2):1-38.
- Blackmore, Susan (1999) *The Meme Machine*. Oxford: Oxford University Press.
- Boyd, Robert, and Peter J. Richerson (2000) "Memes: Universal acid or a better mouse trap?" In *Darwinizing Culture: The Status of Memetics as a Science*, Robert Aunger (ed.), pp.143-162. Oxford: Oxford University Press.
- Dawkins, Richard (1976) *The Selfish Gene*. Oxford: Oxford University Press.
- Dawkins, Richard (1982) *The Extended Phenotype*. Oxford: Oxford University Press.
- Dawkins, Richard (1999) Introduction to Blackmore (1999).
- Dawkins, Richard, and John R. Krebs (1979) "Animal signals: Information or manipulation?" in *Behavioral Ecology*, John R. Krebs and N. B. Davies (eds.) London: Blackwell, pp. 282-309.
- Freeman, Walter J. (1999) *How Brains Make Up Their Minds*. London: Weidenfeld and Nicolson.
- Gil-White, Francisco J. (2001) "L'evolution culturelle a-t-elle des règles?" *La recherche Hors Série* No. 5(Avril): 92-97.
- Henrich, Joseph, and Robert Boyd (in press) "On modeling cognition and culture: Why cultural evolution does not require replication of representations," *Journal of Cognition and Culture*.
- Lewontin, Richard C. (1970) "The units of selection," *Annual Review of Ecology and Systematics* 1:1-18.
- McCrone, John. (1999) *Going Inside: A Tour Round a Single Moment of Consciousness*. London: Faber and Faber.
- Miller, Geoffrey (1999) "Sexual selection for cultural displays," in *The Evolution of Culture*, Robin Dunbar, Chris Knight, and Camilla Power (eds.). Edinburgh: Edinburgh University Press.
- Pinker, Steven (1997) *How the Mind Works*. London: Penguin.
- Sperber, Dan (2000) "An objection to the memetic approach to culture," in *Darwinizing Culture: The Status of Memetics as a Science*, Robert Aunger (ed.), pp. 163-74, Oxford: Oxford University Press.
- Sperber, Dan (1996) *Explaining Culture: a Naturalistic Approach*. Cambridge, MA: Blackwell.
- Sperber, Dan, and Deidre Wilson (1986) *Relevance: Communication and Cognition*. Cambridge, MA: Harvard University Press.
- Tooby, John and Leda Cosmides (1992) "The psychological foundations of culture," in *The Adapted Mind*, Jerome H. Barkow, Leda Cosmides and John Tooby (eds.), pp. 19-136, Oxford: Oxford University Press.