DAN SPERBER

SELECTION AND ATTRACTION IN CULTURAL EVOLUTION*

Suppose we give ourselves the goal of developing mechanistic and naturalistic causal explanations of cultural phenomena. (I don't believe, by the way, that causal explanations are the only ones worth having; interpretive explanations, which are standard in anthropology, are better at answering some of our interrogations.) A causal explanation is mechanistic when it analyses a complex causal relationship as an articulation of more elementary causal relationships. It is naturalistic to the extent that there is good ground to assume that these more elementary relationships could themselves be further analysed mechanistically down to some level of description where their natural character would be wholly unproblematic.

The kind of naturalism I have in mind aims at bridging gaps between the sciences, not at universal reduction. Some important generalisations are likely to be missed when causal relationships are not accounted for in terms of lower level mechanisms. Other valuable generalisations would be lost if we paid attention only to lower level mechanisms. If we want bridges, it is so as to be able to move both ways.

Social sciences explanations are sometimes mechanistic, but they are hardly ever naturalistic (with a few exceptions in demography or in historical linguistics). They fail to be naturalistic if only because they freely attribute causal powers to entities such as institutions or ideologies the material mode of existence of which is left wholly mysterious. If we want to develop a naturalistic program in the social sciences, we must exert some ontological restraint and invoke only entities the causal powers of which can be understood in naturalistic terms.

Here is a proposal: let us recognise only human organisms in their material environment (whether natural or artificial), and focus on these organisms' individual mental states and processes, and on the physical-environmental causes and effects of these mental things.[1] And here is how, having so restricted our ontology, we might approach the social. A human population is

*I am grateful to Ned Block, Gloria Origgi, John Maynard Smith, and Eliott Sober for useful comments on earlier versions of this paper.

[1]Of course, it is contentious that mental things can be naturalised. If they cannot, if no bridges can be thrown between the psychological and the neurological levels, then this reduces to a proposal to bridge the social with the psychological and the ecological sciences. If, as I believe, a naturalistic program in psychology is well under way, then this is a proposal to bridge the social and the natural sciences through psychology and ecology.

M.L. Dalla Chiara et al. (eds.), Structures and Norms in Science, 409–426.

inhabited by a much wider population of mental representations, i.e., objects in the mind/brain of individuals such as beliefs, fantasies, desires, fears, intentions, etc. The common physical environment of a population is furnished with the public productions of its members. By public production I mean any perceptible modification of the environment brought about by human behaviour. Productions include bodily movements and outcomes of such movements. Some productions are long lasting, like clothes or buildings, other are ephemeral, like a grin or the sounds of speech.

Typically, public productions have mental representations among their causes and among their effects. Mental representations caused by public productions can in turn cause further public productions, that can cause further mental representations, and so forth. There are thus complex causal chains where mental representations and public productions alternate. Given that public productions are likely to have many mental representations among their causes and conversely — a point to which I will revert —, the image of a causal lattice or a causal network might be more appropriate. I will keep using the standard phrase "causal chain," it being understood that every link in this chain may be attached to many, both up and down the causal path.

Of particular interest are causal chains from mental representations to public productions to mental representations etc., where the causal descendants of a representation resemble it in content. The smallest ordinary such causal chain is an act of successful communication. Typically, the public productions that are involved in communication are *public representations* such as linguistic utterances. Public representations are artefacts the function of which is to ensure a similarity of content between one of their mental causes in the communicator and one of their mental effects in the audience.

Communication is one of the two main mechanisms of transmission, imitation being the other. Transmission is a process that may be intentional or unintentional, co-operative, or non-co-operative, and which brings about a similarity of content between a mental representation in one individual and its causal descendant in another individual. Most mental representations are never transmitted. Most transmissions are a one-time local affair. However, it may happen that the recipient of an act of transmission becomes a transmitter in turn, and the next recipient also, and so on, thus producing a long chain of transmission, and a strain of mental representations (together with public representations in cases of communication) linked both causally, and by similarity of content. Fast moving rumours, and slow moving traditions are paradigmatic examples of such cultural causal chains.

When you have a strain of representations similar enough in content to be seen as versions of one another, it is possible and often useful to produce yet another public version to represent in a prototypical manner their partly common content. Thus we talk of *the* belief in metempsychosis, *the* recipe of the Yorkshire pudding, or *the* story of King Arthur, each identified by a content. These are, of course, abstractions, at least as much so as *the* zebra,

the Doric order, or *the* Russian peasant. It is tempting to see all the concrete representations that can be identified by means of a prototypical version as having the same content, with only negligible variations, as imperfect replicas of one another, but replicas nevertheless. Once this is done, it is but one step to seeing all tokens of the "same" representation as forming a distinct class of objects in the world, just as all zebras are commonly seen as forming a natural kind. Granting such unity to strains of representations makes it possible to use, in order to develop a causal explanation of culture, one of intellectual history's most powerful tool: the Darwinian idea of selection.

On this approach, cultural representations are self-replicating representations. They replicate by causing those who hold them to produce public behaviours that cause others to hold them too. Occasionally representations "mutate", possibly starting a new strain. The task of explaining the contents and evolution of a given culture can be seen, then, as one of finding out which representations are more successful at replicating, under what conditions, and why. Versions of this idea have been defended by, among others, Donald Campbell (1974), Richard Dawkins (1976, 1982), Cavalli-Sforza & Feldman (1981), Lumsden & Wilson (1981), Boyd & Richerson (1985), and William Durham (1991).[2] Dawkins has coined the name "memes" for cultural replicators and the success of this word has been such that it could be seen as confirming, or at least illustrating, the very idea of a meme. I will focus this discussion on Dawkins' memes, espoused in philosophy by Daniel Dennett (1991, 1995), and developed in anthropology by William Durham (1991). My argument extends unproblematically (mutatis mutandis, of course) to all the other proposals mentioned.[3]

Since Sperber (1985b -see also Sperber 1996), I have been arguing that there is a severe flaw in attempting to develop a naturalistic explanation of cultural evolution on the basis of the Darwinian model of selection. I am not moved by any reservation towards Darwinism. Quite the opposite, I believe that Darwinian considerations have a central role to play in the explanation of human culture by helping us answer the fundamental question: what biological, and in particular what brain mechanisms make humans cultural animals, with the kinds of culture they have. In other words, to characterise "human" in the phrase "human culture" we must draw on biology, hence on evolutionary theory, hence on the Darwinian model of selection. It is "culture" in the phrase "human culture" that calls for some different and, I believe, novel thinking, remaining however, as we will see, within the broad Darwinian approach.

[2]See Sober 1991.

[3]On the other hand, the view of cultural evolution put forward by Pascal Boyer (1993, chapter 9) and my own are very near. Boyer's arguments and mine are in part similar, and in part complementary. Boyer offers a detailed discussion of the models of Lumsden and Wilson (1981), Boyd and Richerson (1985) and Durham (1991). There is also a good deal of convergence with Tooby and Cosmides (1992). Two other original and important approaches, that of the cognitive anthropologist Ed Hutchins (1995), et that of the philosopher Ruth Millikan (1984, 1993) would deserve a separate discussion.

My two basic points have been 1) that representations don't in general replicate in the process of transmission, they transform, and 2) that they transform as a result of a constructive cognitive process. Replication when it truly occurs is best seen as a limiting case of zero transformation. These remarks of mine have been taken as an emphatic way of making a correct but unimportant point to the effect that replication is not perfect.[4] But after all, hasn't Dawkins himself pointed out that "no copying process is infallible," and that "it is no part of the definition of a replicator that its copies must all be perfect" (Dawkins 1982: 85)?

Dawkins however is aware of the looming problem:

> The copying process is probably much less precise than in the case of genes: there may be a certain 'mutational' element in every copying event [...]. Memes may partially blend with each other in a way that genes do not. New 'mutations' may be 'directed' rather than random with respect to evolutionary trends. [...]. These differences may prove sufficient to render the analogy with genetic natural selection worthless or even positively misleading (Dawkins 1982: 112).

Dawkins main interest and most relevant contribution are to point out that the mechanism of Darwinian selection is by no means limited to biological material but may apply to replicators of any substance and any kind.[5] Computer viruses, for instance, are successful replicators (alas!) of a non-biological kind.

Provided that certain conditions obtain, replicators will undergo Darwinian selection. The two main conditions are that there should be variations among replicators, and that different types of replicators should differ in their chances of being replicated. In the case of the selection of genes, the source of variation is random mutation, that are, actually, failures to replicate properly. For selection to operate on replicators capable of mutating, a further condition has to be fulfilled. It has to do with the rate of mutation. Obviously, if genes mutated, not just occasionally, but all the time, they wouldn't be replicators anymore, and selection would be ineffective. How much mutation is compatible with effective selection? Here is George Williams's answer:

> The essence of the genetical theory of natural selection is a statistical bias in the relative rates of survival of alternatives (genes, individuals, etc.). The effectiveness of such bias in producing adaptation is contingent on the maintenance of certain quantitative relationships among the operative factors. One necessary condition is that the selected entity must have a high degree of permanence and a low rate of endogenous change, relative to the degree of bias (Williams 1966: 22–23)

[4]This is how I seem to be interpreted by Dennett (1995: 357–359).

[5]A theme developed in detail by Ruth Millikan 1984.

In fact, interesting replicators — genes in the biological case — can be characterised as entities replicating well enough to undergo effective selection.[6]

In the case of genes, a typical rate of mutation might be one mutation for one million replications. With such low rates of mutation, even a very small selection bias is enough to have, with time, major cumulative effects.[7] If, on the other hand, in the case of culture there may be, as Dawkins acknowledges, "a certain 'mutational' element in every copying event," then the very possibility of cumulative effects of selection is open to question.

There are, of course, bits of culture that do replicate. Some people copy chain-letters. Medieval monks copied manuscripts. Many traditional artefacts are replicas. Thus a pot may be copied by a potter, and some of her pots may be copied by other potters, and so on for many generations of pots and potters. This slow manual reproduction process has been superseded in Modern times by more and more sophisticated technologies such as printing, broadcasting, or email forwarding, which allow massive replication. However the number of artefactual replicas of a would-be cultural item is only a poor and indirect indicator of its genuine cultural success. Waste paper baskets and their electronic counterparts are filled with massively replicated but unread junk, while some scientific articles read by only a few specialists have changed our cultural world. The cultural importance of a public production is to be measured not by the number of its copies in the environment but by their impact on people's minds.

The most blatant cases of replication are provided by public productions rather than by mental representations. However, when the a public replica is produced by an individual rather than by a machine, this production is caused by an intention or a plan of the individual, i.e. a mental representation. Mental representations causing the production of public replicas can themselves be seen as mental replicas of mental representations. Jane's mental representation of a pot caused her to make a pot in conformity to this representation. This pot was seen by John and caused in him the construction of a mental representation identical to Jane's. John's representation caused him to produce a pot identical to Jane's pot, etc.

The question then arises as to whether the true memes are public productions — pots, texts, songs, etc. — that are both effects and causes of mental representations, or, as Dawkins (1982) argues, mental representations that are both causes and effects of public productions. With either option, however, there are similar problems. To begin with, most cultural items, be they mental

[6]Williams goes as far as to suggest that, "In evolutionary theory, a gene could be defined as any hereditary information for which there is a favourable or unfavourable selection bias equal to several or many times its rate of endogenous change. The prevalence of such stable entities in the heredity of populations is a measure of the importance of natural selection" (Williams 1966: 25).

[7]See Wilson & Bossert, 1971: 61–62.

or public, have a large and variable number of mental or public immediate ascendants.

Leaving aside mechanical or electronic reproduction, cases of new items produced by actually copying one given old item are rare. When you sing Yankee Doodle, you are not trying to reproduce any one past performance of the song, and the chances are that your mental version of the song was the child of the mental versions of several people. Most potters producing quantities of near-identical pots are not actually copying any one pot in particular, and their skill is typically derived from more than a single teacher (although there may be one teacher more important than the others, which complicates matters still).

Generally, if you are serious in describing bits of culture — individual texts, pots, songs or individual abilities to produce them — as replications of earlier bits, then you should be willing to ask about any given token cultural item: Of which previous token is it a direct replica? In most cases, however, you will be forced to conclude that each token is a replica not of one parent token, nor (as in sexual reproduction) of two parent tokens, nor of any fixed number of parent tokens, but of an indefinite number of tokens some of which have played a much greater "parental" role than others.

You might want then to envisage that this process of synthetic replication of a variable number of models is done by a natural equivalent of a morphing program (i.e. a program that takes, say, the image of a cat and that of a man as input, and produces the image of a creature somewhere between the cat and the man as output). Just as in a morphing program, different inputs can be given different weights: you can have your cat-man more like a cat or more like a man, and Jill's skill and her pots may be more like Joan's than like Jane's, though still owing both to Joan's and Jane's skills and pots.

The model that comes to mind now is less immediately reminiscent of Darwin notion of selection than of the notion of "influence" much used in History of Ideas and in Social Psychology. In the case of selection, genes succeed or fail to replicate, and sexual organisms succeed or fail to contribute half the genes of a new organism. Thus descent relationships strictly determine genic similarity (ignoring mutations). Influence, on the other hand, is matter of degree. Two pottery teachers may have shared the same pupils and have therefore the same number of cultural descendants but their common cultural descendants may be much more influenced by one teacher than by the other. The resulting pots too may be descendants of both teachers' pots, but more like the pots of the one than like the pots of the other.

There are nevertheless commonalties between the meme model and the influence model. Both involve an idea of competition. Both define a measure of success, in terms of the number of descendants in one case, in terms of the degree and spread of influence in the other case. Both predict that the most successful items will dominate culture, and that culture will evolve as a result of differences in success among competing items. The meme model might

be seen as a limiting case of the influence model: the case where influence is either 100% or 0%, i.e. where descendants are replica. Actually, formal models of influence in Social Psychology tend to concentrate on this limiting case (e.g. Nowak, Szamrej, & Latane 1990).

Both the meme model and the influence model see human organisms as agents of replication or synthesis, with little or no individual contribution to the process of which they are the locus. At most, the replicating agent may, to some extent, choose what to replicate, and the synthesising agent may choose not only what inputs to synthesise, but also the weights to give to the different inputs. Among the factors of either reproductive or influential success, then, there is the attraction that various possible inputs hold for the agents. However once inputs (and weights in the case of synthesis) have been chosen, the outcome of a successful process of replication or synthesis is wholly determined. Moreover, in these two views, mental representations involved in cultural transmission never contain more information that the inputs they are supposed to represent or synthesise.

The influence model is right, against the meme model, in treating replication in cultural transmission not as the norm but as a limiting case (of 100% influence). Both are wrong however in assuming that, in general, the output of a process of transmission is wholly determined by the inputs (and weights, in the case of influence) accepted or chosen by the receiving organism. The relative stimulus-drivenness involved in both models is not the norm of cultural transmission; it also is a limiting case. Not much of culture is transmitted by means of simple imitation or averaging. Medieval monks copying manuscripts — apparently perfect examples of cultural replication — understood what they copied , and, on occasion, corrected what they took to be a mistake in earlier copying on they basis of what they understood. In general, human brains use all the information they are presented with not to copy or synthesise it, but as pieces of more or less relevant evidence to construct representations of their own.

Let me give three brief illustrations.

Consider your views on President Clinton. They are likely to be very similar to the views of many, and to have been influenced by the views of some. However, it is unlikely that you formed your own views simply by copying, or by averaging over other people's. Rather, you used your own background knowledge and preferences to put into perspective information you were given about Clinton, and to arrive by a mixture of affective reactions and inferences at your present views. The fact that your views are similar to many other people's may be explained, not at all by a copying process, and only partly by an influence process; it may crucially involve the convergence of your affective and cognitive processes with those of many people towards some psychologically attractive type of views in the vast space of possible views on Clinton.

Take languages as a second illustration (see also Boyer 1993: 281). Languages are, at first blush, superb examples of memes: complex skills transmit-

ted from generation to generation and similar enough across individuals to allow communication. However, as Noam Chomsky argued long ago (Chomsky 1972, 1975, 1986), a language such as "English" is an abstraction to which correspond, in speaker's minds, mental grammars, and, in the environment, linguistic utterances. Individuals never encounter other people's grammar, or representations of other people's grammars. Individual learners develop their own grammar on the basis of a large but limited set of linguistic utterances. Different individuals encounter very different sets of utterances. Acquiring a language does not consist in imitating these utterances. In fact, most utterances are never repeated. New utterances are not derived either by averaging over, or recombining old ones.

Clearly, what happens in language acquisition is that utterances are used as evidence for the construction of a mental grammar. How constraining is this evidence? Chomsky argued — quite convincingly — that the linguistic evidence available to the child vastly underdetermines the grammar. Moreover, many utterances, being grammatically defective, are bad evidence for the grammar to be constructed. Given this underdetermination, and given the differences in the inputs available to different children, the fact that children do each develop a grammar, and moreover that, within the same community, these grammars converge, raises a deep problem. Again we owe to Chomsky at least the general form of the solution: There is a domain-specific genetically specified language acquisition device in every child's mind. In the vast space of possible uses of the stimulation provided by linguistic utterances, children are attracted towards their use as evidence for grammar construction. In the vast space of logically possible grammars, they choose among just a few psychological possibilities, and they end up converging towards the one grammar psychologically available in the vicinity of the evidence they have been given. Just as it does not matter on which side of the trough you drop the ball, it will roll to the bottom, it does not matter which French utterances a French child hears, she will construct a French grammar.

As a third example, take Little Red Riding Hood. As good an example of a meme as you will ever get. Here, there is no question that many individuals who hear the tale do aim at retelling it, if not verbatim, at least in a manner faithful to its content. Of course they don't always succeed, and many of the public versions produced by one teller for the sake of one or a few hearers differ from the several standard versions.

For instance, suppose an undergifted teller has the hunters extract Little Red Riding Hood from the Big Bad Wolf's belly, but forgets the grand mother. Meme theorist might want to argue — and I would agree with them — that such a version is less likely to be replicated than the standard one. The meme theorists' explanation would be that this version is less likely to have descendants. This is indeed plausible. There is another explanation, however, which is also plausible: hearers whose knowledge of the story derives from this defective version, are likely to consciously or unconsciously correct the story

when they tell it, and, in their narrative, to bring the grand mother back to life too. In the logical space of possible versions of a tale, some versions have a better form, that is, a form seen as being without missing or superfluous parts, easier to remember, and more attractive. The factors that make for a good form may be rooted in part in universal human psychology, and in part in a local cultural context. In remembering and verbalising the story, tellers are attracted towards the better forms. Both explanations, in terms of selection and in terms of attraction, may be simultaneously true, then; the reason why defective versions have fewer replicas may be both that they have fewer descendants, and that the descendants they have are particularly unlikely to be replicas.

I hope the general idea is now clear: There is much greater slack between descent and similarity in the case of cultural transmission than there is in the biological case. Most cultural descendants are transforms, not replicas. Transformation implies resemblance: the smaller the degree of transformation, the greater the degree of resemblance. However resemblance among cultural items is greater than one would be led to expect by observing actual degrees of transformation in cultural transmission. Resemblance among cultural items is to be explained to some important extent by the fact that transformations tend to be biased in the direction of attractor positions in the space of possibilities.

How, then should cultural transmission be modelled? Isn't a Darwinian selection model still the best approximation, to be corrected maybe, but not discarded? To try and answer, let me review the case by means of simple and sketchy formal considerations.

Imagine a population of items that are individually capable of begetting descendants, and that have a limited life-span. Let us imagine that these items come in 100 types with relationships of similarity among the types such that we may represent the space of possibilities by means of a 10 by 10 matrix. Imagine some initial stage (which might, for instance, have been experimentally contrived) where we have a random distribution of say 10,000 items among the 100 types. Suppose we examine our population after a number of generations, and observe a different distribution. While he overall size of the population is roughly the same and there is still a scatter of items across the space of possibilities, now some types are much better represented than others. More specifically, we observe that items tend to be concentrated at, and around two types (see figure 1). Imagine that repeated observations show this distribution to be roughly stable.

A well-know kind of explanation of such a state of affairs would be that some of the types were initially better at replicating, and increased in numbers until a kind of ecological equilibrium was reached, where the more successful types can keep up a higher representation than others. Suppose however that we investigated the manner in which items in this population actually beget descendants, and discovered that an offspring is *never* of the same type as its parent! Rather, the offspring is always of one of the eight types adjacent

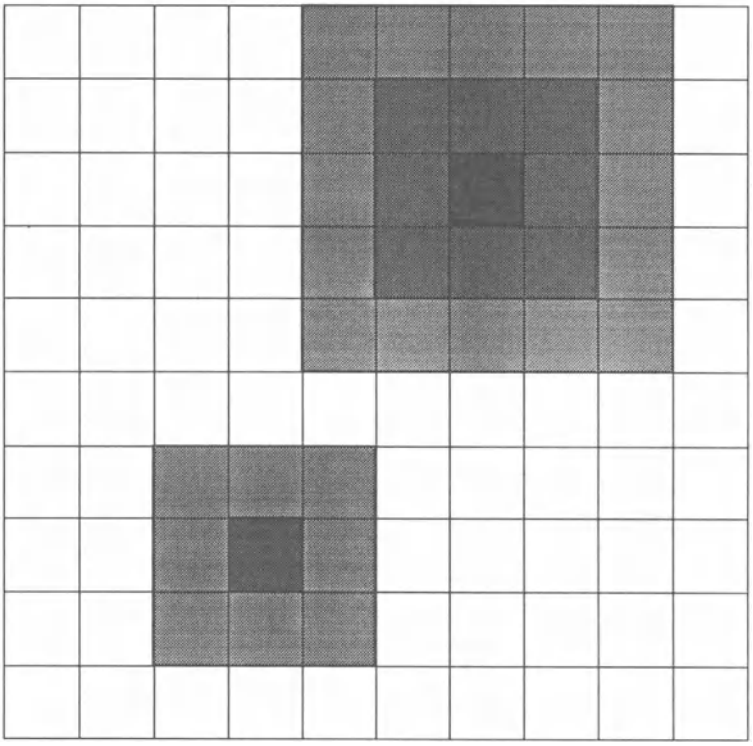

Fig. 1. The space of possibilities: After a few generations, density is greater in the two shaded areas.

in the matrix to that of its parent (fig 2). Ecological equilibrium among differently endowed replicators cannot be the explanation then, since this is a transformative, rather than a reproductive descent system.

An alternative explanation would start from the assumption that the eight possibilities for the offspring of parent of a given of type are not equiprobable. A parent is more likely to beget a transform that differs from it in a given direction. Suppose that the differences in transformation probabilities are such that the matrix has two attractor points.[8] If we trace the descent line of a given item, then, it will not look like a true random walk on the space of possibilities, but will seem, rather, to be attracted towards one or the other of these attractors (figure 3). If the departure point of a descent line is far from the attractors, then it is likely that the arrival point will be near one of them.

[8]Sophisticated notions of attractors ("strange attractors" in particular) have been developed in complex systems dynamics and may well turn out to be of future use in modelling cultural evolution, but a very elementary notion of an attractor will do for the present purpose.

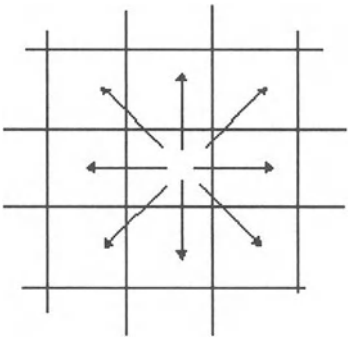

Fig. 2. An item begets descendants of neighbouring types

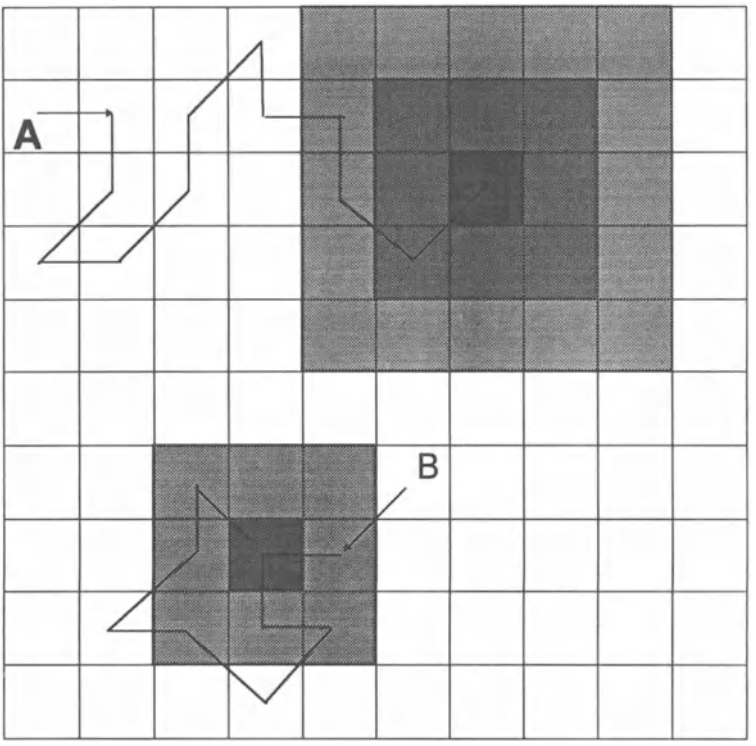

Fig. 3. Descent lines tend to move towards an attractor (line A), or to stay in its vicinity (line B).

It the departure point is near an attractor, then it is likely that the whole line will stay in its vicinity.

If items begot replicas, differences in initial reproductive success and ecological equilibrium would explain the observed distribution. Since items beget transforms, differences in transformation probabilities provide a better explanation.

Transformation and replication can combine. For instance, all the types might have at all times an equal probability of, say, 1/9 of replicating instead of transforming. In this case, although some replication would occur, the difference in distribution between the types would entirely be explained by differences in the probabilities of given transformations. Or the probability of an item replicating rather than transforming might differ according to its type. We could, then, in principle have a dual explanation invoking both reproductive success and attraction. However, in such a case, for the sake of generality and simplicity, replication is better considered as zero transformation, and reproductive success in a given region as defining, or contributing to define, that region as an attractor.

The multiplicity and varying number of "parents," or sources, for the same item, which is, as we noted, a typical aspect of cultural evolution, is also more naturally handled in terms of attraction. The generation of new items in a space of possibilities with attractor regions is to be expected somewhere between existing items and nearby attractors. A simple distance metric is not to be expected, however. The actual mechanisms of generation determine both which items will be transformed, and in which manner.

The attraction model easily incorporates the influence model as a special case: the case where the space of possibility does not exhibit an attractor in the vicinity, and where a simple metric will predict where new items will emerge.

Note that an attractor, as I have characterised it, is an abstract statistical construct, like a mutation rate or a transformation probability. To say that there is an attractor is just to say that, in a given space of possibilities, transformation probabilities form a certain pattern: they tend to be biased so as to favour transformations in the direction of some specific point, and therefore clustering at and around that point. An attractor is not a material thing, it does not physically "attract" anything. To say that there is an attractor is not to give a causal explanation; it is to put in a certain light what is to be causally explained, namely a distribution of items and its evolution, and to suggest the kind of causal explanation to be sought, namely the identification of genuine causal factors that bias micro-transformations.

The existence of attractors is to be explained by two kinds of factors: psychological and ecological. The environment determines the survival and composition of the culture-bearing population; it contains all the inputs to the cognitive systems of the members of the population; it determines when, where and by what medium transmission may occur; it imposes constraints

on the formation and stability of different types of public productions. The mental organisation of individuals determines which available inputs get to be processed, how they are processed, and which information gets to guide behaviours that, in turn, modifies the environment.

Psychological factors interact with ecological factors at several levels corresponding to different time scales: that of biological evolution, that of social and cultural history, that of the cognitive and affective development of individuals, and that of micro-processes of transmission.

It is within the time scale of biological evolution that emerges a species endowed with mental capacities that make cultural transmission possible. The role of biology is not just to make cultures possible, without affecting their character and content. The picture of the human mind/brain as a blank slate on which different cultures freely inscribe each its own worldview, the picture of worldviews as integrated systems wholly determined by socio-cultural history, these pictures, which many social scientists still hold, are incompatible with our current understanding of biology and psychology.

The brain is a complex organ. Its evolution has been determined by the environmental conditions that could enhance or hamper the chances our ancestors had to have offspring throughout phylogeny. There are good reasons to believe that the brain contains many sub-mechanisms or "modules" which evolved as adaptations to these environmental opportunities and challenges (Cosmides & Tooby 1987, 1994, Tooby & Cosmides 1989, 1992). Mental modules, i.e., adaptations to an ancestral environment, are crucial factors in cultural attraction. They tend to fix a lot of cultural content in and around the cognitive domain for the processing of which they are specialised (Sperber 1994).

Evolutionary pressures are likely to have favoured not only the emergence of specialised mental mechanisms, but also some degree of cognitive efficiency within each of these mechanisms and in their mutual articulation. At any given time, humans perceive more phenomena than they are able to pay attention to, and they have more information stored in memory than they can exploit. Cognitive efficiency involves making the right choices in selecting which available new information to attend to, and which available past information to process it with. The right choices in this respect consist in bringing together input and memory information the joint processing of which will provide as much cognitive effect as possible for as little mental effort as possible.

Deirdre Wilson and I have argued that the effect/effort balance in the processing of any given piece of information determines its degree of *relevance* (Sperber & Wilson 1986/1995). We claim that human cognitive processes are geared towards the maximisation of relevance. Most factors of relevance are highly idiosyncratic and have to do with the individual's unique location in time and space. Some factors of relevance however are rooted in genetically determined aspects of human psychology. Thus, the processing of stimuli for which there exist a specialised module is comparatively less effort demanding

and therefore potentially more relevant. For instance, from birth onwards, humans expect relevance from the sounds of speech (an expectation often disappointed, but hardly ever given up).

It is plausible that individuals are equipped so as to tend to optimise the effect/effort balance not just on the input side, but also on the output side. Public productions, from bodily movements, to speech, to buildings, even when they are modelled after some previous productions, are likely to be attracted towards forms where the intended effect can be achieve at a minimal cost.

Human culture has been around long enough for biological evolution to have been affected by it in return. Gene-culture "co-evolution" (Boyd & Richerson 1985, Durham 1991, Lumsden & Wilson 1991) helps explain in particular the existence in humans of abilities that are specifically geared towards cultural interaction such as the language faculty (Pinker & Bloom 1990, Pinker 1994). Gene-culture co-evolution is, however, too slow a process to explain cultural changes in historical time.

Generation after generation humans are born with essentially the same mental potential. They realise this potential in very diverse ways. This, then, is due to the different environments, and in particular to the different cultural environments, in which they are born. However, from day one, an individual's psychology becomes is enriched and made more specific by cultural inputs. Each individual quickly becomes one of the many loci among which is distributed the pool of cultural representations inhabiting the population. The cultural history of a population is both that of its pool of cultural representations, and that of its cultural environment. These macro-ensembles — the pool of representations and the environment — evolve as an effect of micro-processes where the causes belong to the environment and the effect to the pool, or conversely.

In general, the phrase "cultural environment" is used very loosely and refers to an ensemble of meanings, values, techniques, etc.. The cultural environment so understood has little to do with the physical environment. Its ontological status is at best very vague; its causal powers are mysterious. By "cultural environment," I mean an ensemble of material items: all the public productions in the environment that are causes and effects of mental representations. The cultural environment so understood blends seamlessly in the physical environment of which it is a part. The causal powers it exerts on human minds are unproblematic: public productions affect sense organs in the usual, material way. They trigger the construction of mental representations the contents of which are partly determined by the properties of the triggering stimuli, and partly by preexisting mental resources.

Cultural attractors emerge, wane, or move, some rapidly, others slowly, some suddenly, over historical time. Some of these changes have standard ecological causes: overexploited ecological niches lose their economic attraction; rarely walked paths become overgrown; some practices tend to increase, and

others to decrease, the size of the populations that might be attracted to them, etc.

Most historical changes of attractors, however, are to be explained in terms of interactions between ecological and psychological factors of attraction of a kind specific to cultural evolution. The cultural environment causes at every instant the formation of mental representations some of which themselves cause public productions, and so on. This process modifies the relative density of mental representations as well as that of public productions in different areas of the space of possibilities. In particular, density tend to increase in the vicinity of attractors. An increase in the density of public productions in the vicinity of an attractor tends to reinforce the attractor, if only because of the increase in probability that attention will go to these more numerous productions. On the other hand, an increase in the density of mental representations in the vicinity of an attractor may weaken the attractor. The repetition of representations having the same tenor may decrease their relevance and bring individuals either to lose interest in them, or to reinterpret them differently.

Established practices (in matters of dress, food, etiquette, etc.) act as strong attractors. At the same time, because of their expectability, established practices often are low on relevance, while manifest departures from established practices often are an easy way to attract attention and achieve high relevance. Once public productions massively converge towards some cultural attractor, they may foster the emergence of nearby competing attractors. This is illustrated in a dramatic way by the rapid turnover of fashions that quickly lose their power because of their very success.

When, on the contrary, one encounters practices that remain stable for generations, one may suppose that they somehow maintain a sufficient level of relevance in spite of repetition, and try to see whether such is indeed the case, and why. A repetitive practice may remain relevant because its effects are. This is, for instance, the case with technological practices, the economic effects of which are important to peoples welfare or even survival. A repetitive practice may remain relevant because it is in competition with other practices and the choice of one rather than the others by a given individual at a given time may be quite consequential. This is, for instance, the case with practices used to assert one's belonging to some minority. A repetitive practice may remain relevant because different individuals are in competition for the right of engaging in it, and because success in this competition is consequential. This is, for instance, the case with ritual practices marking promotion to some desired status. A repetitive practice may remain relevant because, without hardly modifying its public form, it lends itself to different interpretation according to the agent, the circumstances, the stage in the life cycle. This possibility of renewed re-interpretation is typical of religious practices (see Sperber 1975).

Within the time scale of individual life cycles, ecological and psychological factors also interact in a specific manner. At different stages of their psycho-

logical development, individuals are attracted in different directions. Initially, the main psychological factors of attraction are genetically determined, but experience, that is, the cognitive effects of past interactions with the environment, becomes an increasingly important factor of attraction.

For much of childhood, information that allows the child to develop competencies for which she has an innate disposition is attended to and used for this purpose. The child becomes a competent speaker, a competent climber, thrower, catcher, eater, drinker, a competent manipulator of objects, a competent recogniser of animals, a competent predictor of other people's behaviour, and so on. In all these domains, new information easily achieves relevance because it meets the not yet satisfied needs of specialised modules. As basic competencies are acquired, however, attraction shifts to new information relevant in the context of the already acquired basic knowledge. It shifts in particular to cultural information, for instance religious representations, that seems to challenge basic competencies. It shifts also to information relevant to the various goals that the individual has acquired the ability to conceive and pursue.

The contribution of individuals to cultural transmission varies throughout their life. Not only do they transmit different amounts and different contents, they also transform what they transmit in different directions and transmit to different audiences according to their stage in life. The amplitude of transformations also varies with the age and social role of the communicator and with those of the audience. In some configurations, a relatively conservative communication appears more relevant. In other configurations, the search for relevance demand innovation. From the point of view of individuals, cultural attractors seem to move following a path that, actually, combines historical changes and the individuals' own movement in their life cycle and in their social relationships.

It is the micro-processes of cultural transmission that make possible gene-culture coevolution, and that bring about the historical evolution of culture, and the cultural development of individuals. My argument has been that these micro-processes are not, in general, processes of replication. I am not denying that replications occur and play a role in cultural evolution. I am arguing that replications are better seen as limiting cases of transformations. Constructive cognitive processes are involved both in representing cultural inputs and in producing public outputs. All outputs of individual mental processes are influenced by past inputs. Few outputs are mere copies of past inputs. The neo-Darwinian model of culture is based on an idealisation, which is good scientific practice. However this idealisation is itself based on a serious distortion of the relevant facts, and this is where the problem lies.

The neo-Darwinian model and the ideas of replication and selection seemed to offer an explanation of the existence and evolution of relatively stable cultural contents. How come, if replication is not the norm, that among all the mental representations and public productions that inhabit a human popula-

tion and its common environment, it is so easy to discern stable cultural types such as common views on Bill Clinton, tellings of Little Red Riding Hood, English utterances, and also shake-hands, funerals, and pick-up trucks? For two reasons. Firstly, because, through interpretative mechanisms the mastery of which is part of our social competence, we tend to exaggerate the similarity of cultural tokens and the distinctiveness of types (see Sperber 1985a, chapter 1, 1993). Secondly because, in forming mental representations and public productions, to some extent all humans, and to a greater extent all members of the same population at a time are attracted in the same specific directions.

Even if it contrast with the neo-Darwinian models of culture put forward by Dawkins and others, the model of cultural attraction that I have outlined is, quite obviously, of Darwinian inspiration in the way it explains large-scale regularities as the cumulative effect of micro-processes. The culture of a given population is described as a distribution of mental representations and public productions. Cultural evolution is explained as the cumulative effect of differences in frequency between different possible transformations of representations and of productions in the process of transmission. In the study of cultural evolution, borrowing the Darwin's selection model, is not the only way, and may not be the best, to take advantage of Darwin's most fundamental insight.

CNRS and École Polytechnique
Paris

REFERENCES

Boyer, Pascal (1993). *The naturalness of religious ideas*. Berkeley: University of California Press.

Boyd, Robert & Peter J. Richerson (1985). *Culture and the evolutionary process*. Chicago: The University of Chicago Press.

Cavalli-Sforza, L.L. & M.W. Feldman (1981). *Cultural transmission and evolution: A quantitative approach*. Princeton University Press.

Chomsky, Noam (1972). *Language and mind*. New York: Harcourt Brace Jovanovich.

Chomsky, Noam (1975). *Reflections on language*. New York: Pantheon.

Chomsky, Noam (1986). *Knowledge of language: Its nature, origin, and use*. New York: Praeger.

Cosmides, L. & J. Tooby (1987). From evolution to behavior: Evolutionary psychology as the missing link. In J. Dupré (Ed.) *The latest on the best: Essays on evolution and optimality*. Cambridge Mass.: MIT Press.

Cosmides, L. & J. Tooby (1994). Origins of domain specificity: the evolution of functional organization. In L.A. Hirschfeld & S.A. Gelman (eds), *Mapping the mind: domain specificity in cognition and culture*. New York: Cambridge University Press, 85–116.

Dawkins, Richard (1976). *The selfish gene*. Oxford: Oxford University Press.

Dawkins, Richard (1982). *The extended phenotype*. Oxford: Oxford University Press.

Dennett, Daniel (1991). *Consciousness explained*. Boston: Little, Brown.

Dennett, Daniel (1995) *Darwin's dangerous idea: Evolution and the meaning of life*. New York: Simon & Schuster.

Durham, William H. (1991) *Coevolution: Genes, culture and human diversity*. Stanford: Stanford University Press.

Hutchins, Edwin (1995). *Cognition in the wild*. Cambridge, Mass: MIT Press.

Lumsden Charles J. & E.O. Wilson (1981). *Genes, mind and culture*. Cambridge, Mass: Harvard University Press.

Millikan, Ruth Garrett (1984). *Language , thought, and other biological categories*. Cambridge, Mass: MIT Press.

Millikan, Ruth Garrett (1993). *White Queen psychology and other essays for Alice*. Cambridge, Mass: MIT Press.

Nowak, Andrej, Jacek Szamrej & Bib Latane (1990). From private attitude to public opinion: A dynamic theory of social impact. *Psychological Review. 97*. 362–376.

Pinker, S. (1994). *The language instinct*. New York: Morrow.

Pinker, S. & P. Bloom (1990). Natural language and natural selection. *Behavioral and brain sciences*, 13 (4), 756–758.

Sober, E. (1991). Models of cultural evolution. In P. Griffiths (ed.), *Trees of life: Essays in the philosophy of biology*. Kluwer. 17–38.

Sperber, Dan (1975a). *Rethinking symbolism*. Cambridge: Cambridge University Press.

Sperber, Dan (1985a). *On anthropological knowledge*. Cambridge: Cambridge University Press.

Sperber, Dan (1985b). Anthropology and psychology: towards an epidemiology of representations (The Malinowski Memorial Lecture 1984). *Man* (N.S.) 20, 73–89.

Sperber, Dan (1993). Interpreting and explaining cultural representations. In Gisli Palsson (ed.), *Beyond boundaries: Understanding, translation and anthropological discourse*, Oxford: Berg. 162–183.

Sperber, Dan (1994). The modularity of thought and the epidemiology of representations. In L. A. Hirschfeld & S. A. Gelman (eds), *Mapping the Mind: Domain specificity in cognition and culture*, New York: Cambridge University Press, 39–67.

Sperber, Dan (1996). *Explaining culture: A naturalistic approach*. Oxford: Blackwell

Sperber, Dan & Deirdre Wilson (1986/1995). *Relevance: Communication and cognition*. Oxford: Blackwell. *Second edition* 1995, Oxford: Blackwell.

Tooby, J. & Cosmides L. (1989). Evolutionary psychology and the generation of culture, Part I: theoretical considerations. *Ethology & sociobiology*, 10, 29–49.

Tooby, J. & Cosmides L. (1992). The psychological foundations of culture. In J. Barkow, L. Cosmides & J. Tooby (Eds.) *The adapted mind: Evolutionary psychology and the generation of culture*. New-York: Oxford University Press.

Williams, George C. (1966). *Adaptation and natural selection*. Princeton: Princeton University Press.

Wilson, E. & Bossert W. (1971). A primer of Population biology. Sunderland, Mass: Sinauer.