

WARNING CONCERNING COPYRIGHT RESTRICTIONS

The copyright law of the United States (Title 17, United States Code) governs the making of photocopies or other reproduction of copyrighted material.

Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If electronic transmission of reserve material is used for purposes in excess of what constitutes "fair use", that user may be liable for copyright infringement.

FROM ANIMAL SOCIETIES TO HUMAN SOCIETIES

Despite the obvious similarities between a termite mound and a human city, there are profound differences between the mechanisms that lead to cooperation in the two cases. One important feature of human societies—namely the recognition of individuals—already exists in some social mammals and birds. Although insects may recognize group membership, they do not recognize individuals. In contrast, a monkey recognizes other members of its troop as individuals, and behaves differently towards them. As the phrase 'pecking order' implies, the members of a flock of chickens sort themselves into a linear dominance hierarchy: this probably requires individual recognition. Students of baboons and other monkeys have observed the formation of alliances, in which two or more individuals support one another in conflicts with other members of the group. Such alliances may be based on genetic relatedness, but this is not always so.

The essential points are that, in higher animals, social interactions within a group depend on individual recognition, and that an individual's behaviour towards another depends both on genetic relatedness, and on a memory of previous interactions with that individual.

The characteristics of human societies

It is often said that the defining characteristic of human societies is cultural inheritance; that is, individuals in a society acquire their beliefs and behaviour, their knowledge and skills, by learning from previous generations, and not by genetic inheritance. There is obviously a lot of truth in this idea, particularly when talking of the differences between one individual and another, or between one society and another. At the level of the individual, differences in political opinions are not caused by differences between genes. At the level of societies, the differences between the inhabitants of London today and in the year 1098, or between the inhabitants of London and Beijing today, are caused culturally and not genetically. Having said this, however, there are some reservations that

need to be made. First, there is some cultural inheritance in animals, a fact that is important when thinking about the origins of human culture. Second, the ability of humans to learn, and to build societies dependent on cultural transmission, is genetic: human societies differ from chimpanzee societies because humans and chimps differ genetically. Third, humans learn some things more readily than others: the human mind is not a blank slate upon which experience can write what it will.

Young rats can acquire a preference for a new food by smelling it on the coat of other rats. This is a kind of cultural inheritance: two groups of rats feed on different foods, and the difference is transmitted by learning. The mechanism has been called 'local enhancement': the adults create an environment in which it is easier for the young to learn. This contrasts with 'observational learning', in which one animal watches what another is doing, and copies it. For example, in Britain, where milk is sometimes delivered to the door in bottles with shiny metal tops, great tits learn to tear open the tops to get at the cream. The habit is culturally transmitted, as is shown by the fact that in some areas, at some times, tits do not do it. One might guess that the mechanism of transmission is observational learning: young tits see their elders opening bottles and copy them. But this is probably not so. Most animals seem incapable of observational learning. The explanation of bottle-opening seems to be that young birds, in flocks that have the habit, encounter bottles with torn tops, and so learn that there is cream to be found by tearing them: they do not watch other tits opening bottles and copy them.

But it is hard to believe that all culturally inherited traits in animals depend only on local enhancement. For example, in some areas of Greece, golden eagles feed largely on tortoises. They are unable to break open the shell with their beaks, so a bird picks up a tortoise, flies up to a considerable height, and drops it onto the rocks below, thus breaking the shell. It would be absurd to suggest that in Greece, but nowhere else, this behaviour in eagles is genetically programmed. A young bird could learn by local enhancement that tortoise shells contain meat that is good to eat. But how, other than by copying, could they learn to fly up carrying a tortoise, and drop it? We give a second example, from chimpanzees, below.

The distinction between local enhancement and observational learning is important, because only observational learning can lead to cumulative cultural change, which is the characteristic feature of human history. By observational learning, young individuals can learn from adults, but also, if one individual stumbles on an improved way of doing something, that improvement can be copied. The result is that change can be continuing, rather than occasional, and that an individual can learn, by copying, a skill that it could never have learnt on its own.

It is clear that humans depend on observational learning, reinforced by teaching, including verbal instruction. As the example of golden eagles shows, observational learning is not unknown in animals. The Swiss zoologist Christophe Boesch gives the following example in chimpanzees. Some, but not all, populations of chimpanzees dip sticks into the nests of driver ants, and feed on the ants that crawl up the sticks. The chimpanzees in Gombe National Park in Tanzania use a different technique from those in the Tai National Park in Côte d'Ivoire, and catch about four times as many ants per minute. Local enhancement could explain why some populations dip for ants whereas others do not, but cannot explain why the Tai chimpanzees continue to use an inefficient technique, when there is no reason why they should not adopt a more efficient one. But continued use of an inefficient technique is what we expect if youngsters copy their elders.

If higher animals, at least sometimes, are able to copy their elders, why is it that continuous cultural change does not occur among them, as it does in humans? Thus one population of chimpanzees may differ from another for cultural reasons, but a given population is not continuously acquiring new habits. The likely explanation is that, in humans, the main mechanism whereby culture is transmitted is language. The nature and origin of language are discussed in the next chapter. At the risk of repetition, two conclusions of that discussion will be summarized here: they are the close analogy between genetic and linguistic methods of transmitting information, and the implications of linguistics for the modular nature of the human mind.

Both the genetic and linguistic systems are able to transmit an indefinitely large number of messages by the linear sequence of a small number of distinct units. In genetics, the sequence of four bases enable the specification of a large number of proteins, and these, by their interactions, can specify an indefinitely large number of morphologies. In language, the sequence of some 30 or 40 distinct unit sounds, or phonemes, specify many words, and the arrangement of these words in grammatical sentences can convey an indefinitely large number of meanings.

Richard Dawkins has emphasized this analogy by introducing the concept of a 'meme', the unit of cultural inheritance analogous to a gene. A meme, he argues, is a replicator. If we invent and tell you a limerick, you may tell it to your friends, and they to theirs: a single original entity—the representation of the limerick in my brain—has replicated, as a gene might replicate. Clearly, there is room for selection: if we invent a funny limerick, it is more likely to replicate than if we invent a boring one. Of course, whether a meme will replicate, or fail, depends on the nature of the human mind, and on the cultural milieu (that is, on the other memes present in the population). But the same is true of a gene: its increase depends on the environment and on what other genes are present.

There are, of course, differences. Genes are transmitted from parent to offspring: memes can be transmitted horizontally, or even from offspring to parent. But there is a deeper difference between genes and memes. Genes specify structures or behaviours—that is, phenotypes—during development: in inheritance, the phenotype dies and only the genotype is transmitted. The transmission of memes is quite different. A meme is in effect a phenotype: the analogue of the genotype is the neural structure in the brain that specifies that meme. When I tell you a limerick, it is the phenotype that is transmitted: I do not pass you a piece of my brain. It follows that, in the inheritance of memes but not of genes, acquired characters can be inherited. If I tell you a limerick and you think of an improvement, you can incorporate it before you pass it on. In this sense, cultural inheritance is Lamarckian. For these reasons, one cannot readily apply population genetic theory to cultural inheritance. But the analogy between memes and genes can be suggestive in a qualitative if not in a quantitative sense. Further, although for simplicity we have illustrated the idea of a meme by the example of a limerick, it can refer to more important examples, like a belief in the Trinity, or a knowledge of how to manufacture gunpowder.

A second implication of linguistics is the notion of a modular mind. A study of language and its acquisition suggests that the ability to speak is not an aspect of general intelligence, but is a special competence. As Noam Chomsky has argued, we have a special 'language organ'. The evidence for this view is discussed in the next chapter. It has led to the suggestion that the brain may have other domain-specific competences: in current jargon, that the brain may be 'modular'. We return to this idea later in this chapter.

From ape to human

All Old World monkeys and apes live in social groups, with the single exception of the orang-utan. Figure 12.1 shows a reconstructed phylogeny, or ancestral tree, of these animals, and the nature of their social structure. Old World monkey females remain in the social group in which they were born: males leave their natal group before sexual maturity, and must enter another group to breed. They are said to be 'female kin-bonded'. In chimpanzees, the situation is reversed: males remain in their natal groups, and females move. As shown in the figure, other hominoids vary in their social systems, but none is female kin-bonded. Robert Foley, from whom Fig. 12.1 is borrowed, argues that male kin-bonding originated in the common ancestor of humans and chimpanzees. This is the most parsimonious assumption, given the phylogenetic tree shown, with humans and chimpanzees more closely related than either is to gorillas. If so, male kin-bonding is the ancestral condition for members of the human family—hominids. The social systems of modern humans are so varied that it is hard

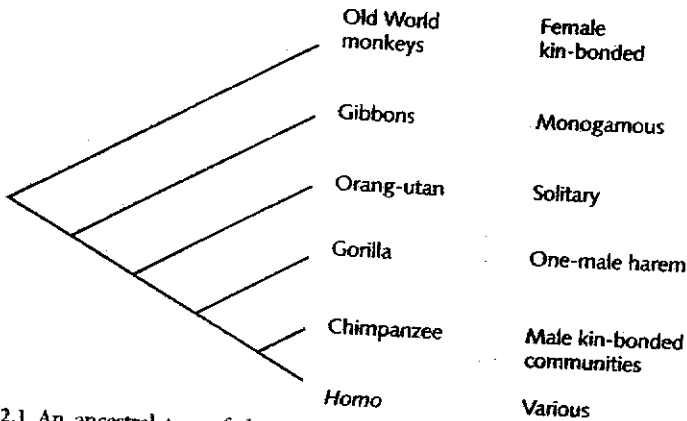


Figure 12.1 An ancestral tree of the primates, with an indication of the social systems associated with different groups, after Robert Foley. The social systems of humans are so varied that it is hard to know what system was characteristic of our hominid ancestors, but Foley suggests that male kin-bonding may have originated in the common ancestor of humans and chimpanzees.

to be sure that this conclusion is correct, but it is the best we can do from the comparative evidence.

The fossil record provides a second source of information about human origins. This is summarized in Fig. 12.2. It is illuminating to compare this record with what is known of human technical achievements, if only for the puzzles that the comparison raises. The australopithecines were bipedal, and lived in wooded grassland. Their relative brain size was only slightly larger than that of apes, and their tool kit was limited and uninventive. In the lineage from *Australopithecus* through *Homo habilis* and *H. erectus* there was a gradual increase in brain size, but rather little technical innovation. The most advanced tool used by *H. erectus* was the handaxe, made from a single block of stone worked on both surfaces, and symmetrical in shape. Such handaxes first appear some 1.4 million years ago, and persist almost unchanged for over a million years: hardly an example of cumulative cultural change.

The most rapid increase in relative brain size has occurred in the past 300 000 years, culminating in the appearance of effectively modern humans some 100 000 years ago. Yet the acceleration in human technical inventiveness, with the appearance of a varied range of tools made of stone, bone, and antler, dates back only 40 000–50 000 years. Burial of the dead, art in the form of cave painting and musical instruments, personal adornment, and trade, originated at much the same time. From about 40 000 years ago, we are faced with evidence of continuing cultural innovation. This raises several problems. Why the delay of 50 000 years between the appearance of the first anatomically modern

humans and the technological revolution? What selective force was responsible for the accelerated increase in brain size 300 000 years ago? When and why did language as we know it originate?

The problems are difficult, because a fossil skull can tell us rather little about the brain that was once inside it, and stone tools little about the society that made them. In *The prehistory of the mind*, published in 1966, British archaeologist Steven Mithen attempted an answer. Although speculative, his book does combine information from palaeontology, archaeology, and psychology to give a convincing answer to these questions. The essence of his argument is as follows. The human mind does indeed contain modules adapted to particular tasks, as suggested by studies of linguistic competence. During much of human evolution, these modules increased in efficiency but remained to a large degree isolated from one another. Language evolved in the first instance to serve social functions, but once grammatical competence had evolved, it provided a means whereby the barriers between modules could be broken down. The burst of creativity during the past 50 000 years resulted from the breaking of these barriers.

Mithen supposes the existence of three mental modules concerned, respectively, with social intelligence, with technical intelligence, and with natural history—that is, with the knowledge of animals and plants needed for efficient foraging. We discuss these in turn.

Social intelligence is a common characteristic of primates. Robin Dunbar, an anthropologist based in Liverpool, has argued that it is the main reason for the increase in brain size in monkeys and apes: his evidence is that there is a striking correlation between brain size in a species, and the size of social groups in that species. A crucial question is how far apes and monkeys have what has been called a 'theory of mind'. To have a theory of mind is to be able to ascribe to others the possession of a mind like one's own, with similar desires and powers of reasoning. There is no convincing evidence that monkeys have such an ability. For example, a vervet monkey gives a different alarm call if it sees an eagle, a snake, or a leopard. But American zoologists Dorothy Cheney and Robert Seyfarth, now at the University of Pennsylvania, who studied these monkeys in the wild, argue that the monkey does not have in mind the knowledge that another monkey may hear its call and respond appropriately; for example, a monkey may continue to call after all others have responded. The behaviour of chimpanzees has convinced most of those who have studied their social behaviour, and skill in manipulation and deceit, that they do indeed have a theory of mind. We can conclude that selection for social intelligence was a major cause of the increase in brain size in monkeys, apes, and humans, and that a theory of mind was already present in the common ancestor of chimpanzees and humans, some 5 million years ago.

Chimpanzees do use tools in the wild; for example, some populations use

stones to crack nuts. Even in captivity, however, their ability to make tools is very rudimentary. Australopithecines used tools, but there is no convincing evidence for deliberate toolmaking. The latter skill first appears associated with the remains of *Homo habilis*, but the tools are little more than irregular chipped stones. *Homo erectus* marks a clear advance, with the manufacture of symmetric handaxes, indicating that the toolmaker had an image of the desired result in mind, and the skill to realize it. But, as mentioned earlier, there remains an astonishing conservatism. Thus there is evidence of a limited increase in technical intelligence, combined with a lack of inventiveness.

There is also evidence for a degree of independence of social and technical intelligence even in modern humans. For example, students of autism have suggested that autistic children are deficient in understanding the behaviour of other humans—what has been called ‘folk psychology’—but better than average in understanding the behaviour of inanimate objects—‘folk physics’.

Finally, there was obviously selection for improved foraging skills, and hence for knowledge of animals and plants, and of their distribution and behaviour. But was this achieved by an increase in general-purpose intelligence, or by the evolution of a specialized natural history module? In favour of the latter, it has been argued that all human societies share certain ideas about the living world. First, that all living things belong to one, and only one, ‘natural kind’. An animal is a dog, or a cat, or a badger, and so on: it must belong to some ‘species’, and cannot belong to two, or to none, and cannot change its species. Second, they share the idea that natural kinds can be classified hierarchically into higher taxa; for example, a dog is a flesh-eater, a mammal (as opposed to a fish, a reptile, etc.), and an animal (not a plant). These universal human attitudes to living things may reflect an innate predisposition. The alternative is that they could be universally believed because they are true, or almost so, and would be learnt by any human society to which a knowledge of the living world was important. A second argument in favour of a special natural history module is the speed with which children acquire these beliefs. But as yet the case for a special module is not decisive.

Thus there is clear evidence for an increase in social, technical, and natural history intelligence in humans, and some hints, particularly from child psychology, of innate and to a degree independent mental modules underlying these skills, analogous to the module responsible for linguistic competence, for which the evidence is much stronger. Mithen points out that, if such modules exist, the extreme conservatism of human prehistory becomes easier to understand. If technical and natural history modules were separate, it would help to explain many features of Lower Palaeolithic tools; for example, the failure to use bone, antler, or ivory in toolmaking, and the uniformity in time and space of stone spear points, despite the range of different food species, in contrast to the great

variability of such points manufactured in the past 30 000–40 000 years. Similarly, if technical and social modules were separate, it would explain the total absence of any form of art, or of personal ornament.

The argument, then, is that the increase in human brain size prior to the emergence of modern humans some 100 000 years ago was associated with an increase in social, technical, and natural history skills, but that these abilities were to a large degree independent. Mithen suggests that the competence for language, including grammar, also evolved in this period, although precise dating is obviously difficult. He ascribes the cultural explosion that began some 50 000 years ago, and that has led to continuous and cumulative cultural change, to a breakdown in the isolation between mental modules caused by the emergence of language. The essential point is that, once words exist for social, technical, and living things, the same grammar can be used to say things about them. For example:

*The stone hit the nut and caused it to break in half.
John talked to Mary and persuaded her to help him.*

In other words, language makes it possible to see the analogies between folk psychology, folk physics, and folk natural history. It is a commonplace that creativity in science and the arts often depends on seeing analogies. If Mithen is right, it was the evolution of language that broke down the barriers between different domains in the mind, and thereby liberated us from the one-million-year conservatism of the Lower Palaeolithic, and made possible the continuous cultural evolution that followed.

Models of human society

Since the time of Plato and Aristotle, philosophers have discussed the nature of society, and how laws can be devised to ensure that it functions harmoniously. It must, therefore, seem somewhat arrogant for two biologists to claim to have something new to say. All the same, we intend to make the attempt. First, we must make it clear what we are *not* going to do. We are not going to argue that human society can be analysed by methods that would be appropriate for an insect society: that is, by explaining the behaviour of individuals in terms of maximizing inclusive fitness (that is, fitness allowing for the effect of an individual's behaviour on the fitness of relatives), and then explaining the structure of the society as a whole as resulting from the interactions of such individuals. The prevalence of cultural inheritance makes such an approach inadequate, although one should not underestimate the role of relatedness in influencing the behaviour of individual humans.

We are, however, going to make use of two ideas derived from biology. The

first is that humans, like other animals, have evolved by natural selection, and therefore have predispositions that ought to make sense in terms of past selection. The second is that complex systems can best be understood by making simple models. It may seem natural to think that, to understand a complex system, one must construct a model incorporating everything that one knows about the system. However sensible this procedure may seem, in biology it has repeatedly turned out to be a sterile exercise. There are two snags with it. The first is that one finishes up with a model so complicated that one cannot understand it: the point of a model is to simplify, not to confuse. The second is that if one constructs a sufficiently complex model one can make it do anything one likes by fiddling with the parameters: a model that can predict anything predicts nothing.

We start, therefore, with a simple model of society, which we call the Social Contract game. The essential assumption is that society consists of a group of equal individuals, behaving rationally. If everyone co-operates, each individual is better off than he or she would be if everyone 'defects'—that is, behaves selfishly. The snag is that, in a society of co-operators, it would pay an individual to defect. If everyone else pays their taxes, it would pay an individual not to do so: but if no one pays taxes, there would be no schools, hospitals, or roads, and everyone would be worse off. Hence, to ensure co-operation there must be an element of compulsion. Suppose, then, that everyone agrees to the following contract, 'I will co-operate: I will join in punishing anyone who defects'. Would not this ensure general co-operation, for the benefit of all? Sadly, it would not. The act of 'joining in punishing' would be costly to the individual, even if not greatly so. The society would hence be invaded by 'free-riders', who co-operated, but did not join in punishing.

The contract can be saved by adding a further commitment, 'I will treat as a defector anyone who does not join in punishing'. It seems that the problem of co-operation is solved: as Immanuel Kant once remarked, it is easy to ensure co-operation between even a race of devils, provided only that they are intelligent. Before considering the weaknesses of the model, we must first ask what qualities individuals must have before they could reach such a contract. First, they must have language. Second, they must have a theory of mind: there would be no point in one individual suggesting to others that they agree to the contract unless he assumed that others had a mind like his own, with similar desires, and similar powers of reason.

The Social Contract game suggests that it should not be all that difficult to ensure co-operation between a group of 'equal individuals, behaving rationally'. The weakness, of course, is the assumption of equality and rationality. We return below to the problem of equality, but first we must query the assumption of rationality. As we write, there are many parts of the world where it is obvious

to outsiders, and to many of the inhabitants, that almost everyone would be better off if they ceased to identify with subgroups—Muslim, Serb, or Croat; Tutsi or Hutu; Jew or Arab; Protestant or Catholic—and worked together for the common good. Yet a sufficient proportion of the population identify with one or other subgroup, rather than with the human population of the region as a whole, to make such co-operation impossible. Why?

The clue is that group identity, and hence behaviour, is influenced by myth and ritual, as well as, and even to the exclusion of, rational self-interest. Historical myths concerning people's origins, reinforced by ritual, are a powerful influence on human behaviour. Why should this be so? What we are seeking is an evolutionary explanation for a universal human characteristic—the ability to be socialized (or indoctrinated, depending on your point of view) by myth. The particular stories, tunes, apparel, and ritual behaviour that bind a group together are clearly cultural, but the capacity to be influenced by them is innate, and calls for an evolutionary explanation. The obvious one is that human groups that could instil group loyalty into their members would be more successful, and so individuals in the group would transmit more of the genes that made group loyalty possible. But is such an explanation plausible?

Again, as so often in this book, we are faced with a conflict between individual and group selection. Archaeological evidence suggests that early human groups were small, which would make between-group selection more plausible. But if there was extensive genetic exchange between groups, as a comparison with other primates suggests is likely, this would greatly reduce the effectiveness of group selection. It seems likely that the capacity for group identification would not have evolved unless it was selectively advantageous within the group, as well as advantageous to the group in competition with others. How could this be so? A plausible answer is that individuals who did not acquire the norms of the group would be penalized by other group members. There is an obvious analogy between this and the nature of the contract required for stable co-operation in the Social Contract game. In that game, co-operation between rational agents requires punishment of those who defect. In the evolutionary scenario we are now suggesting, co-operation is induced by myth and ritual, not by reason, and individual behaviour depends on an innate capacity to be influenced by ritual. We are arguing that such a capacity will evolve by natural selection if two things are true: groups of individuals with such a capacity are more successful, and, within a group, individuals who lack the capacity are punished, just as defectors are punished in the Social Contract game.

Modern societies consist of millions of individuals, not of tens or hundreds. The monotheistic religions are myths inculcating loyalty to the society as a whole, or to all human beings, not just to a small group, although, as the examples given above show, they can all too easily be distorted to fuel hatred

between neighbours. There are subgroups even within societies not divided by language, religion, or history. As we mentioned earlier, a second weakness of the Social Contract game is that it assumes that societies are composed of equal individuals. In agricultural and industrial societies, individuals are not equal. Those who own land, factories, or shops have different options open to them than do peasants, factory workers, or shop assistants. The members of a social class have common interests, and, not unexpectedly, develop myth and ritual to strengthen their struggle to realize those interests: 'The people's flag is deepest red, it's shrouded off our martyred dead'. Curiously, the cohesion of these groups has weakened during the second half of this century, although the inequalities in wealth on which they are based remain. The likely reason is that today's myth-makers, television and the tabloid press, are not controlled by poorer groups within society. There remain groups like the Freemasons bound together by nothing except ritual and self-interest, and the Mafia, in which kin-bonding is also important. Recently, groups based on sex have acquired a political agenda, and are developing myths to strengthen their cohesion.

As in the other transitions described in this book, the emergence of modern society requires the co-operation of entities that, in the past, were independent and competing. Populations of, at the most, a few hundred individuals, with little division of labour except, probably, that between the sexes, have been replaced by societies of many millions, dependent on extensive division of labour. Co-operation depends both on the rational formulation of laws, or social contracts, in the common interest, and on myth and ritual that instil group loyalty. Unhappily, reason can too readily lead to anti-social self-interest, and group loyalty to irrational xenophobia. We need to create legal systems in which self-interest does not lead to social destruction, and myths that extend loyalty to the human species as a whole.

Miroslav Radman, born in ex-Yugoslavia, has recently written of the hatred and cruelty involved in tribal wars between neighbours. We did not see his essay until this chapter had been written. As we do, he seeks an evolutionary explanation for the human instinct that leads to such wars, and suggests that it may lie in the value of human cultural diversity. Although his explanation is not quite the same as ours, he emphasizes, as we do, the importance of myth and ritual in such conflicts, and argues that we need to develop rituals that generate tolerance rather than hatred. Cultural diversity is greatly to be valued, but we need such rituals if we are not to pay a bitter price for it.