

Solving the Puzzle of Human Cooperation

Robert Boyd
Department of Anthropology
University of California
Los Angeles, CA 90095
rboyd@anthro.ucla.edu

Peter J. Richerson
Department of Environmental Science and Policy
University of California
Davis, CA 95616
pjricherson@ucdavis.edu

No consensus on cooperation

Is society an organic whole with each of its many components working together like the organs in a body? Like organisms, societies are composed of many parts which seem to work together enhance their survival. Different people fulfill different, necessary role—subsistence, reproduction, coordination, and defense. Regular exchange of matter and energy guarantees that each component has the resources it needs. Norms, laws and customs regulate virtually every aspect of social interaction, who may marry who, how disputes are resolved, and how verbs should be conjugated. Ritual and religion provide comfort to the sick and fearful, maintain a feeling of solidarity and belonging, and serve to preserve and transmit knowledge through time. Even the simplest human societies seem like complex machines designed for growth and survival.

People have long been divided about whether this metaphor is useful or misleading. Many believe that the appearance of design is real. Functionalism, an old and still influential school in anthropology and sociology, holds that beliefs, behaviors and institutions exist because they promote the healthy functioning of social groups. (Spencer, 1891; Radcliffe-Brown, 1952; Malinowski, 1922; Aberle et al., 1950). Such functionalists believe that the existence of some observed behavior or institution is explained if it can be shown how the behavior or institution contributes to the health or welfare of the social group. The conviction that people are selfish drives others to argue that the appearance of design is an illusion—the complex structure merely reflects a standoff in a struggle among selfish individuals. Such rational individualists, mainly economists, political scientists and philosophers, hold that human choices must be explained in terms of individual benefits; any group benefits are an accidental side effect of selfish individual choices.

This conflict remains unresolved because the competing protagonists espouse irreconcilable views about the causes of human action. Functionalists view people as being shaped by their society. People acquire a belief in the rightness in the norms and customs of their culture as a result of growing up in that society. Believers in the rational individualism see people as choosing how to behave based on their own interest. People are bound by custom only to the extent that it serves these interests. Let us now consider in more detail why this difference in belief about human nature leads to different views about the function of society.

When functionalists do provide a mechanism for the generation or maintenance of group-level adaptations, it is usually in terms of selection among social groups (Turner and Maryanski, 1979). Rappaport (1984) provides an exceptionally clear statement of the idea:

such conceptions as honor, morality, altruism, honesty, valor, righteousness, prestige, gods, heaven, and hell, [make] group selection important among humans. By group selection I mean the selection for and perpetuation of conventions enhancing the persistence of groups, even though these conventions

can be disadvantageous to those individuals whose actions accord most closely with them (p. 401)

Suppose that societies have a number of functional prerequisites without which they have difficulty surviving. Social groups whose culturally transmitted behaviors, beliefs, and institutions do not provide for these prerequisites frequently become extinct, leaving only those societies with functional cultural attributes as survivors. On this argument, selection among social groups causes societies to be adapted to their circumstances in much the same way that selection among individuals causes organisms to be adapted to their environments. We refer to this process as “cultural group selection” because it involves the differential survival and proliferation of culturally variable groups.

Most rational individualists have paid little attention to the group selection argument advanced by functionalists. The reason, we think, is captured by Margaret Thatcher's well-known (famous on the right, infamous on the left) aphorism—“There is no such thing as society.” There can't be selection among societies, because societies do not have intrinsic properties—people make choices in their own interest. If two societies are different, then there must be some difference in conditions that causes individuals to choose differently. Societies cannot replicate themselves because their properties depend on external conditions, not have any transmissible properties. There are a small minority of rational individualists who have taken an interest in selection among groups, and as we will see later in this chapter, their ideas provide one way of bridging the gap between functional and individual level explanations.

Instead rational individualists have devoted considerable ingenuity to devising selfish explanations for the seemingly group beneficial features of society. The first and most famous of these is Adam Smith's invisible hand. Smith showed how the market could regulate the flow of goods and services in a complex society, so that all the things necessary to sustain a city like London are supplied each day without planning. In recent years, scholars have extended this kind of reasoning to many kinds of social interactions not mediated by markets (Becker, 1981, Coleman 1990, Young, 1998) Many of these authors believe that all of the complexity of human cooperation can be explained solely in terms of individual self-interest.

Sociobiology sharpens the puzzle

Placing the question in an evolutionary framework sharpens the puzzle because it supports the rational individualists' assumption that people are selfish, but it casts serious doubt on rational individualists' conclusion that selfish individuals will form complex cooperative societies.

The proposition that human behavior is a product of organic evolution strongly supports the view that people are selfish. Evolutionary theory predicts that any heritable tendency to behave altruistically toward non-relatives will be rapidly eliminated by natural selection. To see why suppose that some individuals in a population have a heritable tendency to help other, unrelated members of their social group at a cost to themselves. For example, suppose some females were motivated by generalized maternal feelings to suckle the orphaned offspring of other females. Such

“compassionate” females would have fewer offspring on average compared to females who lacked this propensity because the compassionate females would have less milk for their own offspring, and all other things being equal, this would reduce their offspring's survival. Thus, each generation there will be fewer copies of the genes that create the motivation to suckle orphans, and eventually, the tendency will disappear.

Selection will favor selfless behavior in only one circumstance: when it is directed toward genetic relatives. To see why, suppose that some females have a heritable tendency to suckle a sister's offspring when they are in need. Since such offspring have a 50% chance of carrying the same genes as the female's own offspring, selection will usually favor such nepotistic motivations if the increase in fitness of the sister's offspring is more than twice the reduction in fitness of the female's own offspring. This reasoning first elaborated by W. D. Hamilton (1964) is supported by an immense body of field and laboratory observation and measurement. It is certainly possible that humans are unusual in some way that caused them to evolve unselfish motives. However, the burden of proof is on people taking this view to show exactly why humans are odd, and in the absence of a clear demonstration of why we are odd, the straightforward prediction of evolutionary biology is that human actions result from selfish or nepotistic motives.

In other species, complex cooperative societies exist only when their members are close relatives. In most animal species cooperation is either limited to very small groups or is absent altogether. Among the few animals that cooperate in large groups are social insects like bees, ants, and termites, and the Naked Mole Rat, a subterranean African rodent. Multicellular plants and many forms of multicellular invertebrates can also be thought of as eusocial societies made up of individual cells. In each of these cases, the cooperating individuals are closely related. The cells in a multicellular organism are typically members of a genetically identical clone, and the individuals in insect and Naked Mole Rat colonies are siblings.

Evolutionary biologists believe that complex cooperative systems are limited to societies of relatives because such systems are vulnerable to self-interested cheating. The many members of an ant colony cannot easily monitor the behavior of all the other members, thus each has the opportunity to cheat on the system. For example, rather than maintain the colony and feeding the queen's offspring, the worker termite can devote time and energy to laying her own fertile eggs. Since the colony has many members, the effect of each on the functioning of the whole is very small, and therefore, each is better off if he or she does cheat. Division of labor creates further opportunities for cheating because it requires exchanges of “goods and services” whose provision is separated in time.

In contrast to the societies of other animals, virtually all human societies are based on the cooperation of large numbers of unrelated people. This is obviously true of modern societies in which complex tasks are managed by enormous bureaucracies like the military, political parties, churches, and corporations. Markets coordinate the activity of millions of people and allow astonishing specialization. It is also true of the human societies that have characterized the human species since first intensive broad spectrum foraging and later agriculture allowed sedentary settlements. Consider, for

example, the societies of highland New Guinea. Here, patrilineally organized groups number from a few hundred to several thousand. These groups have religious, political, and economic specialists, they engage in trade and elaborate ritual exchange with distant groups, and they are able to regularly organize parties numbering several hundred to make war on their neighbors. Even contemporary hunter-gathers who are limited to the least productive parts of the globe have extensive exchange networks and regularly share food and other important goods outside the family. Other animals do none of these things.

Thus we have an evolutionary puzzle. Our Miocene primate ancestors presumably cooperated only in small groups mainly made up of relatives like contemporary non-human primates. Such social behavior was consistent with our understanding of how natural selection shapes behavior. Over the next 5 to 10 million years something happened that caused humans to cooperate in large groups. The puzzle is: What caused this radical divergence from the behavior of other social mammals? Did some unusual evolutionary circumstance cause humans to be less selfish than other creatures? Or, do humans have some unique feature that allows them to better organize complex cooperation among selfish nepotists.

Solutions to the puzzle.

People have proposed five different kinds of solutions to this puzzle:

1. The “heart on your sleeve” hypothesis holds that humans are cooperative because they can truthfully signal cooperative intentions.
2. “Big mistake” hypotheses propose that contemporary human cooperation results from psychological predispositions that were adaptive when humans live in small groups of relatives.
3. Manipulation hypotheses hold that people either tricked or coerced into cooperating in the interests of others.
4. Moralistic reciprocity hypotheses hold that greater human cognitive abilities and human language allow humans to manage larger networks of reciprocity which account for the extent of human cooperation.
5. Cultural group selection hypotheses argue that the importance of culture in determining human behavior causes selection among groups to be more important for humans than for other animals.

These five are not mutually exclusive, and, in fact we believe that the most likely explanation is some combination of the last two hypotheses.

The Heart on Your Sleeve Hypothesis

In his book, *Passions within Reason*, the economist Robert Frank (1988) argues that humans cooperate with non-relatives because people can reliably detect the true

intentions of others. Frank thinks that people have innate dispositions that cause them to be more or less cooperative and that they can signal their disposition to others by their appearance and demeanor. Truthful signals allow cooperators to preferentially interact with other cooperators and this makes possible the evolution of cooperative behavior among non-relatives. Because thinking and feeling are complex physiological processes, Frank argues, it is more costly for a defector to maintain the appearance of a nice guy than it is for a cooperator. The theory of signaling developed in economics and borrowed by biologists (see Gintis 2000 for a discussion) shows that, with this assumption, natural selection may favor honest signals of intent. Thus, Frank's idea is that we wear our hearts, be they good or bad, on our sleeves.

This argument cannot explain why humans are more cooperative than other animals because it applies with equal force to other animals. If humans must wear their hearts on their sleeves, then why not chimps or baboons? True, humans seem smarter than other creatures (although this may mainly be a matter of perspective), but this fact cuts both ways. Being clever may allow humans to better penetrate deceptions, but it may allow them to better perpetrate deceptions as well. In fact, we would use the evidence from other animals to argue against the importance of Frank's mechanism among humans. If the evidence about intentions that one could get by the mechanisms that Frank invokes was as reliable as the information conveyed by kinship, then we should see much cooperation among non-relatives in nature. Since we see very little cooperation among non-relatives, we conclude that signaling provides much poorer information than kinship. Perhaps the reason is that the costs of a defector deceptively signaling that they are altruists is really not very high. If deceptive signalers of altruism can easily reap most of the benefits provided by true altruists without paying the costs of providing altruism, the signal will become useless.

The Big Mistake Hypothesis

Many people have argued that contemporary cooperation results from dispositions that evolved when humans lived in small groups of close relatives (e.g. Alexander 1974, 1987, Hamilton 1975, Tooby and Cosmides 1989). The genus *Homo* is about two million years old, and the species *Homo sapiens* at least 90,000 years old. Until the spread of agriculture beginning about 11,500 years ago humans probably lived in relatively small groups (although how small is a matter of debate). In such a world, it is argued, selection could favor psychological mechanisms that led to unselective altruism towards fellow group members because all the potential recipients were relatives. With the advent of agriculture and later urban life, the size of human groups increased dramatically, and these same predispositions led to altruism toward non-relatives. For example, most people feel sympathy for others. Few can fail to be moved by a photograph of a starving Somali child, its eyes rimmed with flies. Such feelings cause many people to act, to send money, to organize relief, or to choose a less than lucrative career in international medicine or development. Sympathy of this kind is genetically maladaptive in the contemporary world because it leads indiscriminate altruism, but it would have been fine, the argument runs, when all of the potential recipients were relatives. We call this the "big mistake hypothesis" because almost everything in modern life, trade, religion, government, and science is a mistake from the genes' point of view.

The big mistake hypothesis can be extended to explain inter-group enmity (van den Berghe 1981). In many species of non-human primates, kinship ties within groups are stronger than kinship ties among groups and, as a result, neighboring groups compete for territory or other resources. In such an environment, selection might favor a generalized enmity towards non-group members. Today neighboring groups of hunter are typically physically indistinguishable, but often have different dialects, customs, and styles of manufacturing items like arrow points. Thus, it is argued, selection could have favored the rule: “Be nice to people who talk like you, dress like you, and act like you. Be nasty to everyone else.”

The big mistake hypothesis is a cogent story, but we think the evidence from the study of non-human primates tells against it. Many contemporary non-human primates live in small kin based groups very much like those posited for early humans. However, life in such groups does not evidence much non-selective altruism. In all primate species, the members of at least one sex leave their natal group and join another group where their only relatives will be their own children. In most species, it is males who leave, but there are some species such as chimpanzees in which females emigrate. Members of the other sex, usually females, remain in their natal group and live among female relatives. However, when groups are of any size, some of these relatives will be quite close while others will be quite distant. A great deal of evidence suggests that primates are very sensitive to differences in relatedness, directing costly helping behavior mainly toward relatives. Consider baboons for example. A typical baboon group might consist of about 60 animals, about half of them infants and juveniles, 4 or 5 adult male immigrants, and the rest adult females born in the group. Violent conflict among females occurs hourly, and these conflicts sometimes result in serious injuries. Many conflicts involve coalitions, usually made up of matrilineal kin. As a result, female baboons are organized into a dominance hierarchy in which all members of a matriline occupy adjacent positions. Violent conflict among males is also a common occurrence, but coalitions are much rarer (even absent in the chacma baboon), and more ephemeral. When the dominant male confronts a subordinate, the subordinate will often grab an infant to use as a hostage to defend itself. When a dominant male is displaced by a new immigrant, the new dominant will sometimes kill infants sired in the group before he arrived. Such infanticides increase the new dominant's reproductive success because lactation suppresses ovulation, and killing infants causes their mothers to resume ovulation. Other species would tell a similar story—living primates are very good at discriminating between relatives and non-relatives and behave very differently toward each. It is hard to see why early hominids should have been less discriminating in their behavior.

Evidence from human behavior also tells against the big mistake hypothesis. Humans do generally know who their their real relatives are. No doubt propensities to cooperate with kin are deeply ingrained in human psychology. One excellent body of evidence comes from the seemingly tangential literature on incest avoidance. Westermarck (1894) suggested that there is an innate avoidance of inbreeding. If so, humans must have an innate kin recognition system. The operation of this device is nicely illustrated by the rarity of marriage among Israeli kibbutz age-mates (Shepher 1983) and the poor success of Taiwanese minor marriages (Wolf 1970). In these

famous examples, potential husbands and wives are raised in close companionship as children, much as siblings normally are, even though unrelated in fact. Co-resident age-mates apparently have an innate algorithm that invokes a mating avoidance mechanism. The kin recognition system fails in the rather unusual circumstances discovered by Shepherd and Wolf, but in normal families will function properly as an incest avoidance mechanism. If it functions for mate choice it ought also to function to regulate cooperation. Indeed, much human cooperation is nepotistic and nepotistic inclinations generate much conflict larger scale loyalties.

Manipulation hypotheses

When confronted with an example of what seems like cooperation among non-relatives, some sociobiologists argue the apparent behaviors are really the result of coercion by others. Consider the example of delayed marriage in age graded societies in Africa such as the Maasai of Kenya and Tanzania. Amongst the Maasai, every few years all of the boys near puberty are circumcised. All adult men who are circumcised together belong to an "age set." Immediately after circumcision, a man and all the members of his age set spend a period wandering throughout the countryside in an attempt to prove their bravery, for example by killing a lion on foot armed with only a spear. After several years, the members of the age set graduate to become *morani* or warriors. *Morani* provided the manpower for the successful Maasai military machine that was in the process of conquering most of East Africa when the British arrived at the end of the 19th century. Then after eight or ten years, the age set graduates and become older men who do not fight and are eligible to marry. At first blush the behavior of *morani* seems to be an example of altruism toward unrelated individuals: young, vigorous men seem to postpone marriage in order to benefit all the members of their tribe, most of whom are only very distantly related, if at all.

Harpending, et al. (1987) argue that this is not an example of altruism because the *morani* are forced to fight and to postpone reproduction by the older men in *their* own interest. While the older men may not be as powerful fighters as the younger men, according to Harpending et al., their ownership of resources and their networks of political alliances allow them to dominate the younger men. Older men use their power to force the *morani* to do the fighting and keep the women for themselves. Thus, there is no altruism to explain.

However, this explanation is at best incomplete, and at worst, incoherent. Let us suppose that young men obey older ones because if they do not, the older men will punish them. This fact explains why the young men fight and refrain from marriage, but it is not a complete explanation because it does not explain why older men coerce. Coercion is costly to the individuals doing the coercing. If a young man takes a wife, then some of the greybeards will have to lean on him, and this will consume lots of time, effort, and political capital (as anybody who has had to administer academic discipline knows). Moreover, if this is likely to make the young man and his relatives mad, and motivate them to seek revenge. Warriors in their prime should normally be able to beat greybeards if it comes to physical coercion. Thus, the older men must receive some benefit to compensate them for these costs. Some sociobiologists seem to argue that the fact that the *morani* defend the old men and their delayed marriage allows the old men to marry the young women is sufficient compensation. However,

the fact that the older men benefit *as a class* is irrelevant to individual advantage. All that matters is the individual benefit that results from a older man's *own* coercive acts. He must marry the woman, or benefit himself directly in some other way. Moreover, he will be indifferent about marriages that do not affect him directly. In fact, this explanation is just a different group functional explanation which substitutes benefit to a class for benefit to the whole group.

A manipulation argument is cogent if one argues that each older man enforces late marriage on the few young men with whom he interacts. But then it is incomplete because it does not explain why all old men should have such a homogeneity of interests. It is much more likely that the interests of old men are heterogeneous—some would benefit by denying marriage opportunities to their younger colleagues but some would not because they have as many wives as they can afford, or because they want to acquire a rich or well connected husband for their daughter, or because their own sons want a wife. An individual advantage argument predicts a variety of outcomes that depend on the details of local conditions. Sometimes there are good reasons for homogeneity. Consider why in most hunter gatherer societies no women hunt. The usual argument is that women don't hunt because hunting is incompatible with child care. However, this argument does not apply to some women, for example those who are sterile. Why don't they hunt? Hillard Kaplan argues that the reason is that learning complex subsistence skills like gathering or hunting requires many years of practice. Young boys invest in learning to hunt and young girls don't because on average hunting pays for men, but not for women. By the time a woman discovers she is sterile it is too late to acquire the necessary skill to be a successful hunter. It may be possible to rescue the manipulation hypotheses with an analogous arguments, however it remains to be done.

The manipulation hypothesis exemplifies what might be called the naive cynicism of some sociobiological reasoning. Sociobiologists reject the theoretically unmotivated group-functionalism that bedevils much of social science, and seek to explain seemingly altruistic behavior in terms of individual advantage. While we do favor alternative hypotheses to individual advantage, we also think such cynicism about human motives is an appropriate starting point, a good null hypothesis. However, some sociobiologists seem to accept any cynical explanation as if it were based on individual advantage whereas many if not most cynical explanations of human behavior envision one *group* exploiting another.. The fact that the older men benefit *as a class* is irrelevant to individual advantage, it is a group advantage. To reject the argument that young men defer marriage to benefit the group in favor of the argument that old men coerce younger ones in the their own group interest is to substitute one group functional argument for another. In a similar vein, sociobiologists have argued that restrictions on women's action in pastoral societies exist because they are in men's interest (Strassman, 1991) and that commoners have extensive incest restrictions because it is in the interest of the aristocracy (Thornhill, 1990). These explanations sound like an individual advantage arguments because they invoke a conflicts of interest, but none of them really are. They each invoke group advantage—the advantage to men as a group and to the aristocracy as a group.

As we will see in the next section, it is possible to save the manipulation hypothesis, but only if it is stripped of any sense of exploitation. We will see that the mechanism

that makes manipulation possible is reciprocity, and in its simplest formulation, people, in essence, manipulate themselves in their own group interest.

Moralistic Reciprocity Hypotheses

A number of authors have suggested that human cooperation is based on reciprocity (e.g. Trivers 1971, Wilson 1975, Alexander 1987, Binmore 1994), and that our more sophisticated mental skills allow us to manage larger social networks than other creatures. Two kinds of evidence support this hypothesis. First, reciprocity clearly does play an important role in contemporary human societies all over the world. Second, some measures of brain size are correlated with social complexity—animal species which have small social networks tend to have smaller brains (corrected for body size) than do animal species with large social networks (Dunbar, 1992). The fact that humans have very large brains for their body size suggests that humans can maintain reciprocal relationships in larger groups than other animals. Field and laboratory experiments suggest that monkeys are much smarter about social problems than non-social problems. For example, vervet monkeys do not seem to know that python tracks (which are obvious and unmistakable) predict the presence of pythons, but they do know that their aggression toward another vervet predicts aggression by that individual's relatives toward them (Cheney and Seyfarth 1990) which suggests that solving social problems is an important for brain evolution.

The defining feature of reciprocity is that ongoing interactions allow people to monitor each other's behavior and thereby reward cooperators and punish noncooperators. Beyond this property, there is little agreement among biologists or anthropologists about the details of how reciprocity works. In the simplest models punishment takes the form of withdrawal of further cooperation (for example, Axelrod and Hamilton 1980): I will keep helping you as long as you keep helping me, but if you cheat, I won't help you any more. We will refer to such strategies as “simple reciprocity”. Other authors (e.g. Binmore 1994) argue that punishment takes other forms—non-cooperators are punished by various forms of social ostracism, reduced status, fewer friends, and fewer mating opportunities. Following Trivers (1971) we will call this “moralistic reciprocity.” While these different types of reciprocity are often lumped together, they have very different evolutionary properties.

It is very unlikely that large scale human cooperation is supported by simple reciprocity. There is strong theoretical support for the idea that lengthy interactions between *pairs* of individuals are likely to lead to the evolution of this kind of reciprocating strategy (See Axelrod and Dion 1989, Nowak and Sigmund 1993 for review), but recent work suggests that simple reciprocity cannot support cooperation in larger groups (Boyd and Richerson 1988, 1989). Increasing group size places simple reciprocating strategies on the horns of a dilemma. Strategies which tolerate a substantial number of defectors in the group allow defectors to go unpunished and therefore cannot persist when common because such defectors get the benefits of long term cooperation without paying the cost. Thus, reciprocators must be provoked to defect by the presence of even a few defectors. However, such intolerant strategies cannot increase when rare unless there is a substantial chance that the groups made up mainly of cooperators will form when cooperators are rare and they are extremely sensitive to the existence of errors or uncertainty. This dilemma is not serious when

pairs of individuals interact; very minor perturbations allow reciprocating strategies to increase when rare. As groups become larger, however, both of these requirements become impossible to satisfy.

This conclusion makes intuitive sense. We know from everyday experience that reciprocity plays an important role in friendship, marriage, and other dyadic relationships. We will stop inviting friends over to dinner if they never reciprocate, we become annoyed at our spouse if he does not take his turn watching the children, and refuse to return to the auto repair shop when they do a bad job. However, it is not plausible that each one of a thousand union members stay out on strike because they are afraid that their defection will break the strike. Nor does each member of a Mae Enga war party maintain his position in the line of battle because he fears that his desertion will precipitate wholesale retreat.

Moralistic reciprocity provides a much more plausible mechanism for the maintenance of large scale cooperation. Reciprocators can punish non-cooperators in many ways besides withholding their own cooperation.. Strike breakers can be physically attacked or their property can be vandalized. Even more plausibly they can be socially ostracized—scabs lose status in their community and with it many important benefits of social life. Much the same goes for cowards and deserters who may be attacked by their erstwhile compatriots and shunned by their society, made the targets of gossip, or denied access to territories or mates.

Moralistic reciprocity enforced by such of punishment is more effective in supporting large scale cooperation than simple reciprocity for two reasons: When a simple reciprocator stops cooperating in order to punish defectors, he or she induces other reciprocators to stop cooperating. These defections induce still more defections. Innocent cooperators are in effect punished as much as guilty defectors when the only recourse to defection is to stop cooperating. In contrast, other forms of retribution can be targeted so that only defectors are affected. This means that defectors can be penalized without generating a cascade of defection. Second, with simple reciprocity, the severity of the sanction is limited by an individual's effect on the whole group, which becomes diluted as group size increases. Moralistic sanctions can be much more costly to defectors, and therefore, allow rare cooperators to induce others to cooperate in large groups.

However, there is also a problem with moralistic reciprocity that remains to be explained: why should individuals punish? Remember the problem with the manipulation hypothesis: If punishing is costly and the benefits of cooperation flow to the group as a whole, administering punishment is a costly group beneficial act, and therefore, selfish individuals will cooperate but do not punish. The elder Maasai who punishes a younger man, suffers a cost to himself and provides a benefit to his fellow oldsters. Similarly, the striker who attacks a scab may be injured himself. The striker who shuns a scab may forgo a satisfying friendship, a beneficial business relationship, or even a desirable marriage partner. Thus, as long as the effect of the punishment administered by a single individual will have little effect on the success of the strike, selfish individuals will not punish.

This problem is solved if moralistic reciprocators also punish people who do not punish when they should. This means that moralistic strategies punish non-cooperators, individuals who do not punish non-cooperators, and individuals who do not punish non-punishers. When such strategies are common, rare non-cooperators suffer because they are punished. Individuals who cooperate but do not punish suffer against because they are also punished when they don't punish. In this way, it can pay to punish, even though the cooperation that results is not sufficient to compensate individual punishers for the cost of punishing. The Maasai gray beard who arranges his son to marry will lose status in the community, the striker who fails to shun the scab will herself be shunned, and the Enga warrior who fails to participate in punishing a coward may become a "garbage man" himself. If moralists are common, the cost of punishing rare noncooperators may be small, and cost of being punished may be large, so that even quite costly group beneficial behaviors can be maintained by this mechanism. (See Boyd and Richerson 1992 for a more formal version of this argument. There is also a closely analogous result in economics. See Hirshleifer and Rasmusen 1989, or Binmore 1994.)

There is also a big problem with the moralistic reciprocity hypothesis: It explains how costly group beneficial behavior can persist, but it provides no explanation for why group oriented behavior is more common than *any* arbitrary behavior. Moralistic punishment can stabilize *any* arbitrary behavior—wearing a tie, being kind to animals, or eating the brains of dead relatives. It does not matter whether or not the behavior produces group benefits. All that matters is that when moralistic punishers are common it is more costly to be punished than to perform the sanctioned behavior, whatever it might be. When any behavior can persist at a stable equilibrium, then the fact that reciprocity is a stable equilibrium does not tell us whether it is a likely outcome or not.

What we need to know is: Does the fact that reciprocity leads to beneficial outcomes for everyone make it more likely that moralistic punishment will support reciprocity than, say, eating the brains of dead relatives. If the answer to this question is no, then moralistic reciprocity is not a complete answer to the puzzle of human cooperation. To answer this question we need to know how a population changes through time. Our approach is to model behavioral change as a process of cultural evolution by biased transmission. People differ in their belief and values about what behaviors should be subject to moralistic punishment. Then depending on the composition of the population, some beliefs will yield higher payoffs than others, and therefore people with those beliefs will be imitated, and they will spread. This process will then go on until the population reaches a stable equilibrium. We have analyzed a model of the evolution of moralistic punishment that incorporates these ideas, and it suggests that the existence of a group benefit does not increase the ability of a strategy to spread when it is rare (Boyd and Richerson, 1992). Cultural evolution by biased transmission is equally likely to reach an equilibrium at which people are punished for not eating the brains of dead relatives as an equilibrium at which they are punished for not defending their group.

One answer to this objection (Binmore 1994) is that people don't have to slowly evolve toward an beneficial state through a myopic process of cultural evolution. They can deliberately choose what behaviors should be punished either by getting

together and talking about it, or by delegating their choice to a king, council of elders, or some other authority. There is little doubt that collective choice may account for some cases, but it seems to us that there are many more cases which it cannot account for. Think about wearing a tie. In the business world you will surely be punished if you do not wear a tie, yet convention was never deliberately chosen. Or think about Nuer marriage practices. At the time of the Nuer expansion, there were over 100,000 Nuer living in 12 politically independent, and often hostile tribes. Their marriage rules were not chosen, the cultural group was much, much bigger than any social or political entity that could conceivably do any choosing. Outside modern industrialized nation states, it is quite common for the scale of social interaction and political organization to be much smaller than the scale of cultural similarity, and in such cases, we do not see how deliberation can work. Furthermore, deliberation, when it works, is likely founded on a basis of trust in kings, councils, or voters to create rules in the common interest. At least, modern democratic systems seem to work much better when founded on popular communal institutions and poorly in communities where such civic institutions are weak (Putnam, 1993). Suspicious, rational, selfish punishers may in principle deliberately choose rules to mutually enforce, but we suspect that such a regime will prove empirically if also not theoretically fragile.

Cultural Group Selection

The idea that there is selection among societies has a certain empirical plausibility. The Nuer and the Dinka provide the model. According to Kelly (1985) differences in marriage practices led to the Nuer military superiority over the Dinka. The Nuer have a higher and less flexible bride price than the Dinka, requiring households to maintain close relations with a larger group of allied households so as to be able to assemble the requisite number of cows. The Nuer do not seem to have deliberately invented the practice of requiring a larger bride price for the purpose of gaining military advantage, but the larger web of social connections maintained for brideprice reasons allows them to assemble considerably larger fighting forces than the Dinka and thus for a persistent military advantage. If we except Kelly's analysis, then it follows that selection between the Nuer and the Dinka favored Nuer marriage practices. Cultural group selection requires that:

1. there must be differences among groups,
2. these differences must affect persistence or proliferation of groups, and
3. these differences must be transmitted through time.

If these three conditions hold, then, all other things equal, cultural attributes that enhance the persistence or proliferation of social groups will tend to spread. The Nuer and the Dinka clearly differed, this difference allowed the Nuer to expand at the expense of the Dinka, and new Nuer groups continued to behave like Nuer even though they occupied lands previously held by the Dinka. Although the Nuer example is especially unambiguous, it is not unique. There are many examples of human groups disappearing because they could not cope with the environment or

other human groups, and it is plausible in many cases that some groups survived while others did not because of cultural differences.

Why group selection normally doesn't work.

The big problem with cultural group selection is theoretical: How can cultural variation among groups be maintained against the corrosive effect of self-interest. To understand this problem, consider the following simple model: Suppose that individuals live in groups embedded in a larger population. The groups represent the “societies” whose extinction drives cultural group selection. Suppose, for simplicity, that each individual in the population is characterized by one of two cultural variants. One variant causes individuals to place a high value on group goals compared to personal gain. Under the right circumstances, these “cooperators” will act in the group interest even if it is personally costly. The other cultural variant causes individuals to place a low value on group goals compared to personal gain. In the same circumstances, these “noncooperators” will not cooperate. These cultural variants evolve under the influence of four evolutionary forces:

Biased cultural transmission. People observe the consequences of the behavior of others. If, as we assume, cooperation is costly to the individual, most people will see that noncooperators do not suffer as a result of their noncooperation, and will cease their costly cooperation. However, a smaller number will err and conclude (incorrectly) that cooperation is individually beneficial. Most of the time, biased cultural transmission will cause the proportion of cooperators in each group to decrease.

Random cultural drift. If subpopulations are small there may be random changes due to sampling variation. For example, sometimes by chance most people will erroneously conclude that cooperation does pay and during that generation the frequency of cooperation will increase. This effect will be particularly important if a few prominent people play a disproportionate role in cultural transmission.

Migration among groups. No society is completely isolated from other societies. People leave and are replaced by immigrants who bring with them different culturally acquired beliefs and values. People also borrow ideas from people in other societies. Both kinds of mixing cause the proportion of cooperators in each subpopulation to approach the average proportion in the population as a whole. In other words, migration will cause groups to become more alike.

Selection among groups. The probability that a group becomes extinct is negatively related to the proportion of cooperators—the more cooperators, the lower the probability of extinction. When a group goes extinct its territory is occupied by people from surviving groups.

Each of these forces have different effects. Driven by self-interest, biased transmission causes noncooperative behavior to increase within every group. Group selection works in the opposite direction increasing cooperative behavior. Drift and migration affect the pattern of cultural variation within and among groups. We want

to know, under what conditions will group selection cause the group beneficial behavior to come to dominate the population?

The answer is: Hardly ever. The model we have just sketched is closely analogous to “interdemic group selection” models studied by population geneticists. The only real difference is that biased transmission rather than natural selection leads to the spread of selfish behavior within subpopulations. Extensive analysis of such models by population geneticists indicates that populations will eventually become composed almost completely of noncooperators unless subpopulations are extremely small, there is little mixing among groups, and it is very hard for individuals to discern the costs and benefits of alternative beliefs. Selection among groups fails because biased transmission and migration combine to create a powerful force reducing variation among groups. Migration among groups insures that noncooperators are always present in every group, and biased transmission causes their frequency to increase. In contrast, only force that generates differences among groups—random fluctuations due to small population size and small numbers of individuals colonizing empty habitats—is weak unless groups are very small. Thus there will be little variation among groups even if group selection is very strong, and as group selection can have little effect. The existence of cultural transmission is, by itself, no guarantee that there will be enough variation among groups to allow cooperation will evolve.

The impatient anthropologist might say: “Who cares. We anthropologists know that there is heritable variation among societies, so the group selection will work. If sociobiologists and economists and like-minded folks need an account based on individual decisions, let them worry about it. We should get on with the business of determining the function of cultural practices.” Obviously we think that it is important to reconcile anthropological and evolutionary explanations, but many anthropologists do not. Even so, it is a serious mistake to ignore the problem of the maintenance of variation for cooperation between societies. All societies have internal conflict and the balance between cooperation and selfishness varies considerable from one society to the next. The basic empirical facts point to a dynamic balance between selfish and prosocial motives and much complex structure regarding what behavior is appropriate toward co-members of different groups (e.g., co-members of one’s family, extended family, ethnic group, religion, and nation). Humans are not unconditional cooperators so we cannot assume that any given social institution is group function even if many do have such elements. The cases are quite evidently diverse and the devil is in the details of the evolutionary processes in operation (for a nice comparative study see Knauft, 1993). In the remainder of this chapter we will outline two mechanisms that maintain variation in costly-group beneficial behavior. Each mechanism is consistent with human behavior having been shaped by natural selection. However, each mechanism also has empirical requirements that need to be tested with anthropological data, each mechanism makes additional predictions, and each mechanism points to interesting new questions about the logic of the group selection process.

How group selection can work.

For group selection to be an important process that can generate adaptations at the group level there must be some mechanism that can maintain variation among groups, and not just for some behaviors, but for any arbitrary behaviors. We think that there are at least two such mechanisms: Moralistic reciprocity and frequency dependent bias. Lets see how they work.

Variation is maintained by moralistic reciprocity.

Group selection can be an important process generating group beneficial behavior when group beneficial behavior is enforced by moralistic punishment. Moralistic punishment allows wide variety of individually costly memes to persist once they become common, and as a result moralistic punishment provides a mechanism for preserving variation among groups, variation that then can be subjected to selection at the group level. By itself, moralistic punishment does not provide an explanation of human cooperation, but the combination of moralistic punishment and group selection does.

To show how moralistic reciprocity allows group selection to work, we modify the group selection model described above: Once again there is a population subdivided into a fixed number of groups linked by migration and there are two memes. However, now each meme specifies a different set of behaviors that are to be enforced by moralistic punishment. One specifies that a person's primary loyalty is to his or her kin and people who violate this norm are punished; the second, specifies a person's primary loyalty is to his or her group, and, again, violators are punished. We will call these the "nepotistic" and the "nationalistic" memes, respectively. All other things being equal, people are biased in favor of nepotism, but when the nationalistic group loyalty meme is sufficiently common the effects of punishment overcome this bias and people tend to choose the group loyalty meme. Until recently everybody in the population was nepotistic, but then most people in one group adopted the nationalistic meme. We will return the question of why this should occur, but for the moment we simply regard it as a random event, the group level analog of a mutation. The mutant group has a competitive advantage because nationalistic beliefs allow the organization of larger, and more inclusive corporate activities, and as a result it is more likely to survive as a group, and more likely to grow in size. When a group goes extinct, its territory is occupied by others. The question is: Under what circumstances will this group level advantage cause nationalistic beliefs common in the mutant group to spread throughout the population?

First, biased transmission must be strong enough to keep the nationalism meme common in the mutant group. To see why, suppose this meme is rarely expressed, and rarely punished so it is difficult for immigrants to learn that nepotism will get them in trouble, and therefore biased transmission is weak. As a result, only a few immigrants will have switched to nationalism by the time that more immigrants arrive. This reduces the penalty for being nepotistic, and thus even fewer of these immigrants switch during the next time period, and eventually nepotism becomes advantageous, and the initially nationalistic group becomes nepotistic like the population as a whole. Unless biased transmission is strong enough to resist the

homogenizing effects of migration, nationalism cannot spread by group selection because it cannot persist long enough for selection among groups to occur. In contrast, suppose people with nepotistic beliefs are frequently punished in the nationalistic group so that so that new immigrants whose beliefs differ from the majority rapidly learn that such beliefs get them in trouble and adopt the prevailing norm. Then when yet more nepotistic immigrants arrive they find themselves to be in the minority, quickly learn the local norms, and so maintain the nationalistic meme, that originally arose merely by chance, at high frequency long enough for group selection to have some effect.

Second, group reproduction must preserve cultural differences. Suppose that one of groups in which nepotistic memes predominate goes extinct. If their territory is occupied by individuals drawn from all other groups, the majority of the colonists will be nepotistic, and as a result the new group will be nepotistic. Even if the group in which the nationalistic meme is common never goes extinction, nationalism will not spread because the nationalistic group cannot reproduce itself. In contrast, suppose that new groups are formed by the fissioning of existing groups. Then, most of the time, when a nepotistic group fails it will be replaced by a new nepotistic group, but occasionally it will be nationalistic group that fissions, and when this occurs the nationalistic meme spreads. If nationalistic groups are more likely to survive or more likely to grow and fission the result will be that nationalism will spread.

The fact that this mechanism only works when biased transmission is strong explains why this it is likely to be more important for cultural evolution than for genetic evolution, and, therefore, provides a potential explanation for why humans are more cooperative than other animals. Evolutionary biologists normally think of selection as being weak, and although there are many exceptions to this rule, it is a useful generalization. So, for example, if one genotype had a 5% selection advantage over the alternative genotype this would be thought to be strong selection. So, suppose that a novel, group beneficial genotype has arisen, and that it has become common in one local group where it has a 5% advantage over the genotype that predominates in the population as a whole. For group selection to be important, the novel type must remain common long enough to spread by group selection which is possible only if the migration rate per generation is less than 5% (see Boyd and Richerson, 1990 for details.) But this is not very much migration. The migration rate between neighboring primate groups is on the order of 50% per generation. [Find a survey on migration rates] While migration rates are notoriously difficult to measure, it seems likely that migration rates are typically high among small local groups that suffer frequent extinction. Migration rates between larger, deme sized groups is much lower, but so to will the extinction rate. If one assumes that migration rates are of the same order of magnitude for cultural evolution, then it seems likely that biased transmission can easily maintain intergroup variation.

Variation is maintained by conformist social learning

Frequency dependent bias can also maintain variation among groups. We have shown (Boyd and Richerson 1985, Henrich and Boyd 1999) that natural selection can favor a psychological propensity to imitate the common type. When it is difficult, but not impossible, to determine what is the best way to behave in local environment, you

may increase your chances of doing the right thing by doing what everybody else is doing. The propensity to imitate the common type creates an evolutionary force that causes common memes to become more common and rare memes to become more rare. If this effect is strong compared to the effect of direct bias, then variation among groups will be maintained.

Once again consider our basic group selection model. As before there are a number of groups linked by migration. Now, however, assume that the two memes affect religious beliefs: “believers” are convinced that moral people are rewarded after death while the wicked suffer horrible punishment for eternity, while “non-believers” do not believe in any afterlife. Because they fear the consequences, believers behave better than nonbelievers, they are more honest, charitable, and selfless, and as a result, groups in which believers are common are more successful than groups in which nonbelievers are common. People's decision to adopt one meme or the other is only weakly affected by direct bias. People seek comfort, pleasure, and leisure and this causes them to prefer to behave wickedly, at least on occasion. However, it also causes them to prefer to avoid the endless searing pain of an eternity of burial in a burning tomb. Since people are uncertain about the existence of an afterlife, they are not strongly biased in favor of one meme or the other. As a result they are strongly influenced by the meme that is common in their society. People who grow up surrounded by believers, choose to believe, while those who grow up among worldly atheists do not.

The group-beneficial belief in an afterlife will spread if two conditions are satisfied: Frequency dependent bias must be strong enough to overcome the effect of migration, and new groups have to be formed through the fissioning of existing groups. These conditions are similar to the ones under which moralistic reciprocity allows the spread of group beneficial traits. The reason is that both processes lead to bias against rare memes—moralistic reciprocity because carriers of rare memes are punished, and frequency dependent bias because, averaged over many traits, rareness indicates undesirability. In either case, when the bias against rare types is strong compared to mixing, biased transmission can maintain differences between groups. However, the same bias against rare types means that successful beliefs can spread only if groups reproduce themselves as coherent entities so that successful memes remain common.

The fact that variation is maintained by bias against rare memes also determines what is a group extinction. In genetic models, group extinction requires the death of group members. However, the “extinction” of a cultural group requires only the disruption of the group as a social unit and the dispersal of its members. As long as social extinctions are not too common, such dispersal has much the same effect as physical extinction, because dispersing individuals will be a small fraction of any group they join. In most cases they will temporarily change the frequency of alternative meme in their new group a small amount after which it will return to its original composition (which is determined by the balance of migration and the bias against rare types.)

The form of group selection outlined above can be a potent force even if groups are usually very large. For a group beneficial meme to spread, it must become common in a initial subpopulation. The rate at which this will occur through random drift-like processes (Cavalli-Sforza and Feldman, 1981) will be slow for sizable groups (Lande,

1986). However, this need only occur once. Thus, even if groups are usually large, occasional bottlenecks in some groups could allow group selection to get started. Similarly, environmental variation in even a few subpopulations could provide the initial impetus for group selection.

The difference between moralistic reciprocity and biased social learning is illustrated by the different answers they give to the question: Why do most people growing up in a devout Christian society end up, believing in the tenets of the Christian faith? People who believe that cultural variation is maintained by moralistic punishment would explain this difference in terms of reward and punishment: people who do not adopt Christian beliefs in a devout Christian society are punished, and people who do not punish such heretics (say by continuing to associate with them) are themselves punished. People adopt the prevalent belief because they calculate (not necessarily consciously) that it yields the highest payoff. People who believe that cultural variation is maintained by conformist transmission and similar cultural mechanisms would view young people adopting the tenets of Christianity as accurate descriptions of the world (God exists; He intervenes in human affairs, and so on) because such beliefs are widely held and difficult for individuals to prove or disprove.

Each mechanism has limitations. Moralistic reciprocity can only explain adherence to behaviors that are enforced by punishment. We think that there are many examples of people adhering to costly norms even though there is little risk of punishment, and we think most readers do too—even those who think they don't. Here's a self-test. Do you believe anything that you read in ethnographies? If you do, then you agree with us. Anthropologists go off by themselves to study some exotic culture, and after spending a few years observing and talking to the locals, they return and write up their results. Ethnographers have a big incentive to come back with interesting or impressive results—their careers depend on it—and thus they have a strong incentive to fudge or make up their data. What's more they could cheat with little risk of being caught or punished. Observations are rarely replicated, and even when they are the world is noisy and changeable. Thus, if you think that people only obey norms when they fear of punishment you should distrust all ethnography. While we think that the ethnographic method is far from perfect, we also think that ethnographers for the most part do their best to accurately report their experiences.

Frequency dependent bias can potentiate group selection only if it is strong compared to direct bias, and this can occur only if it is difficult for individuals to evaluate the costs and benefits of alternative memes. Clearly there are many belief group beneficial beliefs whose costs are all too obvious—should you cheat on your taxes, or fake illness to avoid military service? To explain the evolution of honest taxpayers or patriotic volunteers you perhaps need only to invoke norms enforced by punishment. However, there are also many such memes whose affects are hard to judge. Will children turn out better if they are sternly disciplined or lovingly indulged? Is smoking marijuana harmful to one's health? Is academic life a promising career in 2000? These are difficult questions to answer, even with all of the information gathering and processing resources of a late twentieth century industrial state. For most people at most times and most places even questions like, Does drinking dirty water cause disease? and, Can people affect the weather by appeals to the supernatural? are very difficult to answer. While it may be very difficult to determine which meme is best,

the choice may nonetheless have profound effect on people's behavior, including behavior driven by rational choice.

Is the model empirically plausible?

Knowing that cultural group selection is logically possible is useful. Some authors have claimed that social and cultural institutions cannot be understood as benefiting the group because this is inconsistent with a theory built up from the actions of selfish individuals.

Given a showing of logical possibility the interesting question is the empirical one. Have human cultures actually been shaped by cultural group selection?

The models suggest that cultural group selection can lead to the evolution of group beneficial traits if:

1. group disruption and dispersal is common,
2. new groups are usually formed by the fissioning of existing groups, and
3. there is variation among groups, and that this variation affects the ability of a group to survive and produce daughter groups.

If these conditions held true for substantial stretches of human history then it is quite plausible that during such periods that human cultures were shaped by cultural group selection.

To address these questions we (in collaboration with our colleague, Joseph Soltis) have collected the relevant data from the ethnographic literature of Irian Jaya and Papua New Guinea. Until recently, most people in New Guinea lived in small, unstratified societies numbering a few hundred to a few thousand people. Such societies are characteristic of more of human history than the complex, stratified societies that dominate the contemporary world. The highlands of New Guinea were not explored by Europeans until the 1930's and not subjugated by Europeans until after the second World War. As a result, we have high quality ethnographic data describing peoples whose societies had not been drastically altered by a colonial administration. We read as many ethnographies as possible looking for those with accounts of group extinction, new group formation, and variation among local groups (Soltis et al. 1995)

The data from New Guinea indicate that group level extinctions are quite common. You need three numbers to estimate extinction rates: The number of groups in the population, the number of extinctions, and the time period during which the extinctions occurred. We were able to find these data for 5 cultural areas in New Guinea, and all 5 reveal high rates of group extinction. The percent of groups estimated to suffer extinction in each generation (25 years) ranges from about 8% among the Maring to about 31% among the Fore (See the table below). Moreover, more than half of the twenty-eight ethnographies surveyed mention the occurrence of

at least one extinction, and it is likely that this underestimates the actual number because the failure to mention an extinction does not necessarily mean extinctions never occurred.

Region	Number of groups	Number of extinctions	Number of years	% groups extinct every 25 years	source
Mae Enga	14	5	50	17.9%	Meggitt, 1977
Maring	13	1	25	7.7%	Vayda, 1971
Mendi	9	3	50	16.6%	Ryan, 1959
Fore/Usurufa	8--24	1	10	31.2%--10.4%	Berndt, 1962
Tor	26	4	40	9.6%	Oosterwal, 1961

Second, the data indicate that new groups are formed by the fissioning of existing groups. Detailed accounts for two cultures, the Mae Enga and the Mendi, report that new clans are formed when sub-clans grow to the size that allows them to take on the social functions usually performed by clans. Anecdotal evidence from other ethnographies indicates that the same general processes operate in other parts of New Guinea and Irian Jaya as well. We could not find any account in which colonists of new land are drawn from multiple groups, so the group reproduction process would favor the preservation of between-group variation.

Finally, the data also support the existence of cultural variation among local groups, although there is little evidence connecting this variation to local group survival. We have had difficulty finding good ethnographic descriptions of cultural variation among local groups. But for three regions, there is documented local variation. The Ok communities in the fringe highlands vary in terms of ritual and social organization. Tor tribes vary in language and kinship terminology. Abelam groups differentially accepted and rejected non-traditional religions. Furthermore, Barth (1971) argued that group selection has occurred among Ok communities. Ritual variation produced differing levels of centralized organization, and the more centralized groups were able to spread at the expense of the less centralized. At the very least, we could not falsify the hypothesis that cultural variation exists among local groups within larger cultural regions. More generally, if ethnographers are accurate reporters, consequential socio-political differences between different small-scale societies are ubiquitous. (Your experimental data not yet citable?)

Rates of change

The New Guinea data on extinction rates allow us to estimate the maximum rate of cultural change that can result from cultural group selection. For a given group extinction rate, the rate of cultural change depends on the fraction of group extinctions that are the result of heritable cultural differences among groups. If most extinctions are due to non-heritable environmental differences (e.g. some groups have poor land) or bad luck (e.g. some groups are decimated by natural disasters), then group selection will lead to relatively slow change. If most extinctions are due to heritable differences (e.g. some groups have a more effective system of resolving internal disputes), then cultural change will be relatively rapid. The rate of cultural change will also depend

on the number of different, independent cultural characteristics affecting group extinction rates. The more different attributes, the more slowly will any single attribute respond to selection among groups. By assuming that all extinctions result from a single heritable cultural difference between groups, we can calculate the maximum rate of cultural change.

Such an estimate suggests that group selection is unlikely to lead to significant cultural change in less than about 1000 years. The length of time it takes a rare cultural attribute to replace a common cultural attribute is one useful measure of the rate of cultural change. Suppose that initially a favorable trait is common in a fraction q_0 of the groups in a region. Then the number of generations t necessary for it to become common in a fraction q_t of the groups can be estimated using a simple formula derived in Soltis et al. (1995). The time necessary for different parameters is given in following table. If one takes the extinction rate of the Mae Enga as typical, these results suggest that group selection could cause the replacement of one cultural variant by a second, more favorable variant in no less than about 20 generations, or 500 years. Given that not all extinctions result from heritable cultural differences, 1000 years is a reasonable upper bound for the rate of evolution of a single character with a strong influence on group survival.

initial fraction favorable trait	final fraction favorable trait	Extinction rate			
		1.6%	10.4%	17.9%	31%
0.1	0.9	192	40.0	22.3	11.8
0.01	.99	570	83.7	46.6	24.8

This rate of adaptation is not nearly fast enough to explain the many cases of cultural change known to have occurred on much shorter time scales. For example, the introduction of the horse to the Great Plains of North America in the 1500's led to the evolution of the culture complex of the Plains Indians in less than 300 years. If the rates of group extinction estimated for New Guinea are representative of small scale societies, cultures like those of the Great Plains cannot be explained in group-functional terms. There has not been enough time for group selection to drive a single cultural attribute to fixation, even if that attribute had a strong effect on group survival.

This result also suggest that group selection cannot directly justify the practice of interpreting many different aspects of a culture as group beneficial. While both empirical and theoretical uncertainties prevent us from making a quantitative estimate of the rate of evolution of many different cultural attributes, it is clear that group selection will shape many traits more slowly than a single trait. If group selection can cause the substitution of a single trait on 1000 year time scales, the rate for many traits will be substantially longer. We know from linguistic and archaeological evidence that related cultural groups, such as the Nuer and the Dinka, which differ in many cultural attributes often have diverged from a common ancestral group in the last few thousand years. Thus, there has not been enough time for group selection to produce the many group beneficial attributes that distinguish one culture from another.

The result *does* provide justification for interpreting *some* aspects of contemporary cultures in terms of their benefit to the group. The model demonstrates that under the

right conditions group selection can be an important process, and the data from New Guinea suggest that these conditions are empirically realistic. The data also suggest that the rates of group extinction are high enough to cause a small number of traits with substantial effects on group welfare to evolve on time scales that characterize some aspects of cultural change. Group selection cannot explain why all the many details of Enga culture differ from the many details of Maring culture. It might explain the existence of geographically widespread practices that allow large scale social organization in the New Guinea highlands, practices that evolved along with, and perhaps allowed, the transition from band scale societies to the larger scale societies that exist today.

Nor do our data exclude the possibility that cultural group selection explains the increase in scale of sociopolitical organization in human prehistory and history. That evolution is slow even on the millennial scale. Anatomically modern humans appear in the fossil record about 90,000 years ago, yet there is no evidence for symbolically marked boundaries (perhaps indicative of a significant sociopolitical unit encompassing an “ethnic” group of some hundreds to a few thousand individuals) before about 35,000 years ago (Klein 1999). The evolution of simple states from food producing tribal societies took about 5,000 years, and the modern industrial state took another 5,000. Evolutionary processes which lead to change on 100 year time scales cannot explain such slow change unless they are driven by some environmental factor which changes on longer time scales. In contrast, the more or less steadily progressive trajectory of increasing scale of sociopolitical complexity over the past few tens of thousands of years indeed is consistent with adaptation by group selection.

It follows from this picture that much contemporary cultural variation may be functionally significant. Critics of functionalism (Sahlins, 1976; Hallpike, 1986) have argued that the diversity of human societies is incompatible with functional explanations because societies faced with similar environments and using similar technologies exhibit radically different forms of sociopolitical organization. Hallpike (1986) argues that the extent of such differences in simpler societies indicates that the variation is not functional, but neutral, and that it cannot be the result of evolution. The alternative is that they represent variations currently being exposed to a slow process of group selection that is as yet far from equilibrium.

We still have the puzzle of the relatively rapid evolution of seemingly prosocial institutions on much shorter than millennial time scales. Most modern industrial societies have evolved from aristocratic and totalitarian forms of government to representative democracies over the last few centuries. Such societies seem to function better, at least under modern conditions, than their predecessors. None of these societies has gone extinct. One possible explanation for such developments is that over the span of many millennia in the Pleistocene, basic human moral intuitions became adapted to living in group selected cultures. The moralistic punishment mechanism, for example, will penalize individuals whose genes as well as whose memes are not conducive to following social rules. As Aristotle said, humans are those animals that by their nature live in political communities (for him, the tribal-scale Greek polis). Some measure of our prosocial impulses may have become innate, an idea we call the “tribal social instincts” or Aristotle’s hypothesis. (Given some prosocial predispositions and some implementing cultural institutions that engender

trust, people may individually or collectively tend to make decisions in the group interest.

Conclusion

Cooperation in human groups is a serious puzzle because we cooperate extensively with distantly related and even unrelated individuals. Hamilton's (1964) theory of inclusive fitness and Trivers (1971) theory of reciprocal altruism predicted that cooperation can only evolve under highly restricted circumstances. At the time they advanced these ideas, many evolutionary biologists assumed that large-scale animal cooperation was common and that large-scale group selection for altruism was the explanation for it. The work inspired by these theories has resulted in a nearly total triumph for Hamilton and Trivers, although vigorous and sophisticated dissent is still heard (Sober and Wilson, 1998). However, human politics is the only really good example to hand of cooperation on large and even very large scales among non-relatives. For the rest of the organic world Hamilton's and Trivers' mechanisms seem virtually as unexceptional as the law of gravity. Indeed, human societies themselves show ample evidence of individual selfishness, nepotism, and small-scale cabals but nevertheless manage, to highly variable degrees, to provide the benefits of cooperation on considerable scales. While we do not regard the case closed by any means, we think that group selection on cultural variation together with moralistic punishment and tribal social instincts are the strongest current candidates to explain how human societies evolve to fly in a world where strong selective forces favoring only small-scale cooperation act against their existence.

References

- Aberle, D.F., A.K. Cohen, A.K. Davis, M.J. Levy, and F.X. Sutton 1950. The functional prerequisites of a society. *Ethics* 60:100–111.
- Alexander, R. D. 1974. The Evolution of Social Behavior. *Annual Review of Ecology and Systematics*, 5: 325–383.
- Alexander, R. D. 1987. *The evolution of moral systems*. NY: Aldine De Gruyter.
- Axelrod, R., and Hamilton, W.D. 1981. The evolution of cooperation. *Science* 211, 1390-6
- Axelrod, R & Dion, D. 1988. The Further Evolution of Cooperation. *Science*, 242: 1385–1390.
- Barth, F. 1971. Tribes and intertribal relations in the Fly Headwaters. *Oceania* 41:171–191
- Boyd, R. and P. J. Richerson, 1985, *Culture and the Evolutionary Process*, University of Chicago Press, Chicago, IL.

- Boyd, R. and P.J. Richerson, 1988. The Evolution Reciprocity in Sizable Groups. *Journal of Theoretical Biology*, 132: 337–356.
- Boyd, R. and P.J. Richerson, 1989. The Evolution of Indirect Reciprocity. *Social Networks*, 11:213–236.
- Boyd, R. and P. J. Richerson, 1990. Group Selection Among Alternative Evolutionarily Stable Strategies. *Journal of Theoretical Biology*, 145: 331–342.
- Boyd, R. and P. J. Richerson. 1992. Punishment allows the evolution of cooperation (or anything else) in sizable groups. *Ethology and Sociobiology*, 13: 171–195.
- Becker, G. 1981. *A treatise on the family*, Cambridge, Mass. : Harvard University Press,
- Binmore, K. G. 1994. *Game theory and the social contract*, Cambridge, Mass: MIT Press.
- Cavalli-Sforza, L. L. and M. W. Feldman 1981. *Cultural Transmission and Evolution*. Princeton: Princeton University Press
- Cheney D. L. & R. M. Seyfarth, 1990. *How monkeys see the world : inside the mind of another species*, Chicago: University of Chicago Press.
- Coleman, J. 1990. *Foundations of Social Theory*. Cambridge, MA, Havard University Press.
- Dunbar, R.I.M. 1992. Neocortical size as a constraint on group size in primates. *Journal of Human Evolution*. 22: 469–93.
- Frank, R. 1988. *Passions within Reason* New York: Norton
- Hallpike, C.R.1986. *The Principles of Social Evolution*. Oxford: Clarendon Press.
- Hamilton, W.D. 1964. Genetical evolution of social behavior I, II. *Journal of theoretical Biology*, 7: 1–52.
- Hamilton, W. D. 1975. Innate social aptitudes in man: An approach from evolutionary genetics, in *Biosocial anthropology*. Edited by R. Fox, pp. 133–55. London: Malaby.
- Harpending, H. Rogers, A., & Draper, P. 1987. Human Sociobiology, *Yearbook Of Physical Anthropology*, 30: 127-150
- Henrich, J. and R. Boyd. 1998. The Evolution of Conformist Transmission and the Emergence of Between-Group Differences. *Evolution and Human Behavior* 19: 215–242.

- Hirshleifer, R. and Rasmusen, E. 1989, Cooperation in a Repeated Prisoner's Dilemma with Ostracism, *Journal of Economic Behavior and Organization*, 12: 87–106.
- Kelly, Raymond C. 1985. *The Nuer conquest : the structure and development of an expansionist system*, Ann Arbor : University of Michigan Press,.
- Klein, R. G. 1999. *The human career : human biological and cultural origins, 2nd ed.* Chicago: University of Chicago Press.
- Knauft, B.M. 1993. *South Coast New Guinea Cultures: History, Comparison, Dialectic*. Cambridge: Cambridge University Press.
- Lande, R. 1985. Expected time for random genetic drift of a population between stable phenotypic states. *Proceedings of the National Academy of Sciences (U.S.A.)* 82: 7641-7645.
- Malinowski, B. 1922. *Argonauts of the western Pacific*. Prospect Heights, IL. Waveland Press (Reissue 1984).
- Nowak, M. & Sigmund, K. 1993. A strategy of win-stay, lose-shift that outperforms tit-for-tat in the prisoner's dilemma game. *Nature*.364: 56-58.
- Putnam, R.D. 1993. *Making democracy work: Civic traditions in modern Italy*. Princeton: Princeton University Press.
- Radcliffe-Brown, A.R. 1952. *Structure and function in primitive society*. London: Cohen & West.
- Rappaport, R. A. 1984. *Pigs for the Ancestors*. Yale University Press.
- Richerson, P.J. and R. Boyd. 2000. The biology of commitment to groups: the tribal social instincts hypothesis, in *The biology of Commitment*, edited by R.M. Nesse. Russell Sage Foundation. In press.
- Sahlins, M. 1976. *Culture and practical reason*, Chicago: University of Chicago Press
- Sepher, J. 1983. *Incest, the biosocial view*. New York: Academic Press.
- Sober, E., and D.S. Wilson. 1998. *Unto Others: The Evolution and Psychology of Unselfish Behavior*. Cambridge: Harvard University Press.
- Soltis, J., R. Boyd, and P.J. Richerson, 1995. Can group functional behaviors evolve by cultural group selection? An empirical test. *Current Anthropology*. 36: 473–494.
- Spencer, H. 1891. *Essays: Scientific, Political, and Speculative* (3 volumes). London: Williams & Norgate.

- Strassman, B. 1991. The Function of Menstrual Taboos among the Dogon Defense against Cuckoldry? *Human Nature*, 2: 89–131.
- Thornhill, N. W. 1990. Evolutionary significance of incest rules, *Ethology and Sociobiology*, 11: 113-129.
- Tooby, J. & Cosmides, L.. 1989. Evolutionary psychology and the generation of culture, part I: theoretical considerations, *Ethology and Sociobiology*, 10: 29-49.
- Trivers, R.L. 1971. The evolution of reciprocal altruism. *Quarterly Review of Biology*, 46: 35–57.
- Turner, J. H. & A. Maryanski. 1979. *Functionalism*. Menlo Park CA., Benjamin/Cummings.
- van den Berghe, P.L. 1981. *The Ethnic Phenomenon*. New York: Elsevier.
- Westermarck, Edward. 1894. *The History of Human Marriage*. London: Macmillan.
- Wilson, E. O. 1975. *Sociobiology*, Cambridge MA, Harvard University Press.
- Wolf, Arthur P. 1970. “Childhood Association and Sexual Attraction: A Further Test of the Westermarck Hypothesis. *American Anthropologist* 72(3): 503-515.
- Young, H. P. 1998. *Individual strategy and social structure : an evolutionary theory of institutions*. Princeton, N.J. : Princeton University Press,.