Why Does Culture Increase Human Adaptability?

Robert Boyd
Department of Anthropology, University of California, Los Angeles
Peter J. Richerson
Center for Population Biology, University of California, Davis

It is often argued that culture is adaptive because it allows people to acquire useful information without costly learning. In a recent paper Rogers (1989) analyzed a simple mathematical model that showed that this argument is wrong. Here we show that Rogers' result is robust. As long as the only benefit of social learning is that imitators avoid learning costs, social learning does not increase average fitness. However, we also show that social learning can be adaptive if it makes individual learning more accurate or less costly.

KEY WORDS: Social learning; Adaptation; Culture; Sociobiology.

INTRODUCTION

Culture has made the human species a spectacular ecological success. Since the first appearance of tools and other evidences of culture in the archaeological record, the human species has expanded its range from part of Africa to the entire world, increased in numbers by many orders of magnitude, exterminated competitors and prey species, and radically altered the earth's biota.

It is not clear, however, *why* culture improves human adaptability. There has been a lot written about this topic, often in the introductions to papers and books on other topics, but very little careful analysis. In previous work, we (e.g., Boyd and Richerson 1985) suggested that social learning allows us to avoid the costs of individual learning. Learning is costly, and without social learning everybody would have to learn everything for themselves. Teaching, imitation, and other forms of social learning, we argued, allow us to acquire a vast store of useful knowledge without incurring the costs of discovering and testing this

Received April 28, 1994; revised October 31, 1994.

Address reprint requests and correspondence to: Professor Robert Boyd, Department of Anthropology, UCLA, 405 Hilgard Ave., 392 Haines Hall, Los Angeles, CA 90024-1553.

Ethology and Sociobiology 16: 125-143 (1995)

knowledge ourselves. Recently, however, Alan Rogers (1989) has shown that this argument is, at best, incomplete and at worst, plain wrong. Using a mathematical model of the evolution of social learning, he showed that the fact that social learning allows individual organisms to avoid the costs of learning does not increase the ability of that species of organisms to adapt. In fact, in the long run, social learning has no effect at all on the evolving organism's average fitness.

Here we have two goals: First, we argue that Rogers' result is robust, not an artifact of the specific form of his model. To do this, we analyze two models that incorporate Rogers' fundamental assumption that social learning allows individuals to avoid the costs of individual learning, but incorporate quite different assumptions about how social learning works and how the environment varies. Because these models also show that social learning does not increase the average fitness, we conclude that Rogers' result is robust. Culture will not increase the ability of a population to adapt if its only benefit is to allow individuals to avoid learning costs. We then analyze two models of the evolution of social learning that incorporate different assumptions about the evolutionary benefit of social learning. They assume that social learning increases the fitness of individuals who do *not* imitate by reducing the cost or increasing the accuracy of individual learning. In these models, culture does increase the average fitness of populations.

WHY AVOIDING LEARNING COSTS DOES NOT INCREASE AVERAGE FITNESS

Rogers' Model

Rogers' conclusions are based on a mathematical model of the evolution of imitation in a very simple hypothetical organism. These animals live in an environment that can be in one of two states; let's call them wet and dry. The environment has a constant probability of switching from wet to dry each generation, and the same probability of switching from dry to wet, which means that over the long run the environment is equally likely to be in each state. The probability of switching is a measure of the predictability of the environment. When this probability is high, knowing the state of the environment in one generation tells you little about the state of the environment in the next generation. In contrast, when the probability of switching is low, the environment in the next generation is likely to be the same as the environment this generation. There are two behaviors available to the organism: one best in wet conditions and the other in the dry conditions. There also are two genotypes – learners and imitators. Learners figure out whether the current environment is wet or dry and always adopt the appropriate behavior. However, the learning process is costly in that it reduces learners' chances of survival or reproduction. Imitators simply pick a random individual from the population and copy it. Copying does not have any direct

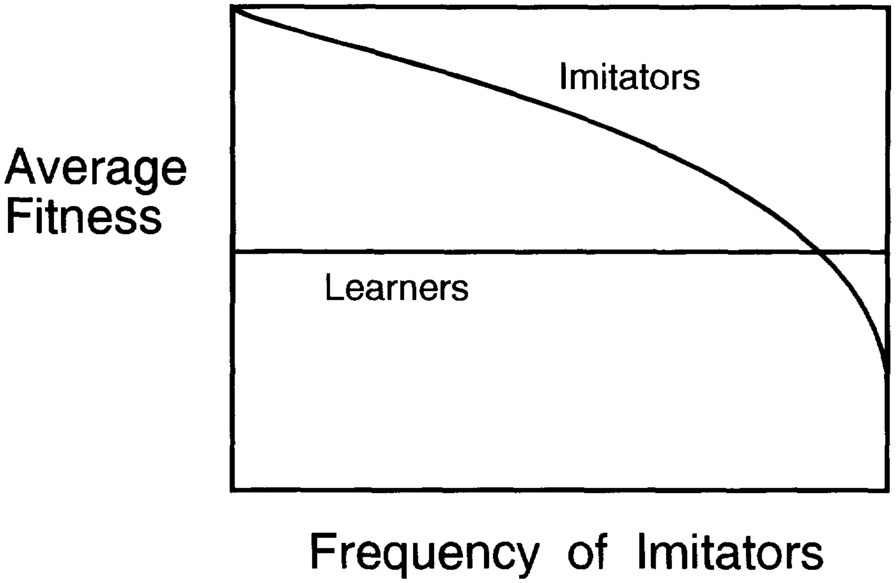

FIGURE 1. The average fitness of learners and imitators as a function of the frequency of imitators in the population. The frequency of learners is one minus the frequency of imitators. This figure is redrawn from Rogers (1989).

effect on survival or reproduction. Rogers then used some simple but clever mathematics to determine which genotype wins in the long run.

The answer is surprising. The long-run outcome of evolution is always a mixture of learners and imitators in which both types have the same fitness as learners in a population in which there are no imitators present. In other words, natural selection favors culture, but culture provides no benefit to the species. The organisms are no better off than they were without any imitation.

To understand the logic of this result, think about the fitness of learners and imitators as the frequency of imitators changes. As is shown in Figure 1, when imitators are rare, they have higher fitness than learners. They are nearly certain to acquire the best behavior because the population is composed of almost all learners, and learners always acquire the right behavior. But imitators don't suffer the cost of learning, so their fitness must be higher than learners. Thus, new mutations that give rise to copying will always be able to invade a population of learners. On the other hand, when learners are rare, they have higher fitness than imitators. When there are very few learners, most of the imitators copy imitators who themselves copied imitators and so on. Because the environment changes periodically, this means that when learners are rare, imitators, in effect, choose behavior at random. In contrast, learners still acquire the best behavior. Thus, rare learners will be able to invade a population of imitators any time that the benefits of learning are sufficient to compensate for its costs. Because both types can increase when they are rare, the population will always be a mix of the two types. But only mixtures in which the two types have the same fitness can be stable long-run outcomes. Since the fitness of the learners is constant, it follows that the evolutionarily stable mix of learners and imitators has the same fitness as a population composed only of learners.

Two Extensions of Rogers' Model

You might think that this paradoxical result is an artifact. After all the model *is* very simple. Perhaps if we add just a little realism, the paradox would go away. But such is not the case. We show below that as long as the only benefit of imitation is the avoidance of learning costs, then changing rules of cultural transmission, the nature of environmental variability, and the number of traits leave Rogers' basic result unchanged.

Spatially varying environment, more than two behaviors, learning errors. Rogers' model assumes that the environment varies in time but not space, that there are only two behaviors, and that learners always acquire the correct behavior. Each of these assumptions can be changed without changing the qualitative result.

Consider a model in which organisms live in an environment that consists of a large number of discrete islands, each with a different environment in which a different behavior is favored by natural selection. The populations on different islands are linked by migration of individuals from each island to all other islands. Thus, in this model the rate of migration measures the predictability of the environment. If migration rates are high, individuals' environments are unlikely to be similar to their parents'. If migration rates are low, most individuals live in environments just like the one their parents lived in. Learners engage in costly learning trials that usually allow them to acquire the locally optimal behavior but also sometimes lead to errors. As is shown in Appendix 1, this model yields the same qualitative result as Rogers' model. Imitation evolves, but does not benefit the population in the long run.

Imitators can detect learners. Unlike the simple organisms in Rogers' model, humans do not blindly imitate a randomly chosen individual. Rather they often evaluate the behavior of many individuals and choose the one that seems best, a process we have labeled *biased transmission* (Boyd and Richerson 1985). Once a beneficial innovation arises, biased transmission allows it to spread through a population without further individual learning. Thus, it seems plausible that if Rogers' model were extended to allow biased transmission, the average fitness of the population might increase. However, a little analysis shows that this intuition is wrong.

Consider a model in which there are learners and imitators. As before, learners always acquire the currently favored behavior but at some cost. After learners learn, each imitator surveys the behavior of n individuals living in its social group. Imitators query each potential model to find out whether it acquired its behavior by copying or by learning. If there is even a single learner in their group, imitators copy the learner and thereby acquire the behavior that is best in the current environment. If there are no learners, imitators copy a randomly chosen individual. This model allows imitators a great deal more information than Rogers' model: They can imitate *n* others rather than one, and they don't copy at random. However, as is shown in Appendix 2, the qualitative result is exactly the same—both types are present, and their long-run average fitness is the same as a pure population of learners.

Why Rogers' result is robust. As Rogers argued in his original paper, his result is robust because it reveals a basic evolutionary property of social learning: The advantage that imitators get from avoiding learning costs cannot increase fitness of a population because the frequency of imitators will increase until this advantage is exactly balanced by the disadvantage that imitators often acquire the wrong behavior. The fundamental logic underlying Rogers' result can be represented graphically as in Figure 1, which plots the expected fitness of learners and imitators as a function of the fraction of imitators in the population. You can see that the fitness of imitators declines as the frequency of imitators increases because the more imitators there are, the more poorly the population tracks the changing environment, the lower the frequency of adaptive behavior, and, therefore, the dumber it is to copy. Moreover, there always have to be some learners in the population, because a population consisting only of imitators behaves at random. Thus, the expected fitness of imitators and learners has to be the same at equilibrium. But the fitness of learners isn't affected by the number of imitators. Thus, at equilibrium the average fitness of the population is the same as a population without culture.

HOW CULTURE CAN INCREASE AVERAGE FITNESS

Thinking about the problem this way points to its solution. Social learning would improve the average fitness of a population if it increased the fitness of *learners* as well as imitators. Consider Figure 2. Here, we assume that the average fitness of learners increases as the frequency of imitators increases, and the paradox disappears—learners and imitators still have the same fitness at equilibrium, but now that fitness is higher than a population composed entirely of learners. Thus, to improve the average fitness of the population, imitation must make individual learning cheaper or more accurate.

Of course, this formal possibility would be of little importance if there were no plausible means by which increasing the amount of imitation would cause individual learning to be more efficient. However, we suggest that there are at least two ways that imitation can benefit learners.

Imitation Allows Selective Learning

Imitation can increase the average fitness of learners by allowing individuals to learn more selectively. Learning opportunities often vary. Sometimes it may be

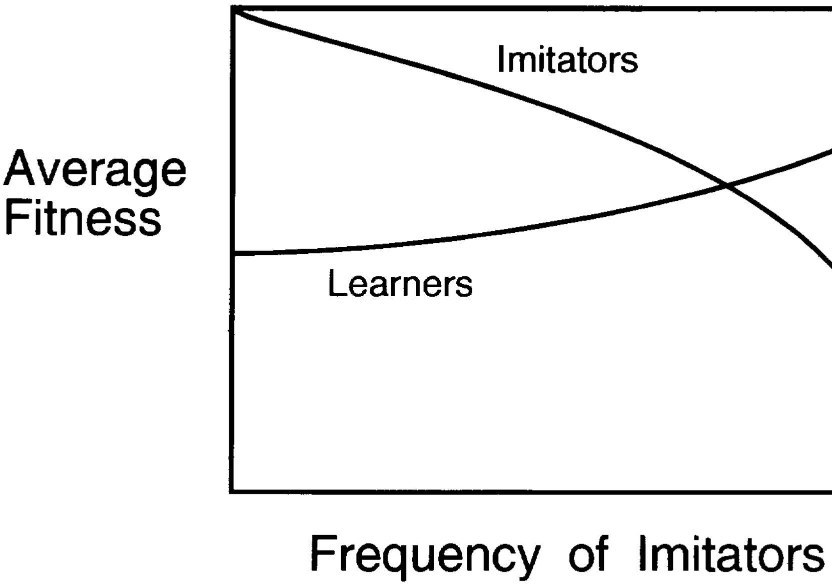

FIGURE 2. If increasing the frequency of imitators reduces the cost or increases the accuracy of individual learning, then the average fitness of the population can be increased by imitation.

easy to determine the best behavior while other times it may be very difficult. Without imitation, an organism must rely on learning even when it is difficult and error prone. In contrast, an imitating organism can learn when learning is cheap and accurate, and imitate when it is costly or inaccurate. The following model shows that imitation plus selective learning can increase average fitness in a population even when most individuals imitate.

As before, consider a population that lives in an environment that switches between two states, and assume that there are two behaviors, one best in each environmental state. However, now suppose that all individuals attempt to discover the best behavior in the current environment. Each individual experiments with both behaviors, and then compares the results. The results of such experiments vary for many reasons, and, thus, the behavior that is best during any particular trial may be inferior over the long run. To avoid errors, individuals adopt a particular behavior only if it appears *sufficiently* better than its alternative. The larger the observed difference in the payoffs between the two behaviors, the more likely it is that the behavior with the higher payoff actually is best. By insisting on a large difference in observed payoff, individuals can reduce the chance that they will mistakenly adopt the inferior behavior. Of course, being selective will also cause more trials to be indecisive, and in that case, they imitate a randomly chosen individual. Thus, there is a tradeoff: You can increase the accuracy of learning, but only by also increasing the probability that learning will be indecisive, and you will have to rely on imitation. The exact nature of the trade-off depends on the probability distribution of the outcome of learning trials. In Appendix 3, we analyze a model in which the observed difference in

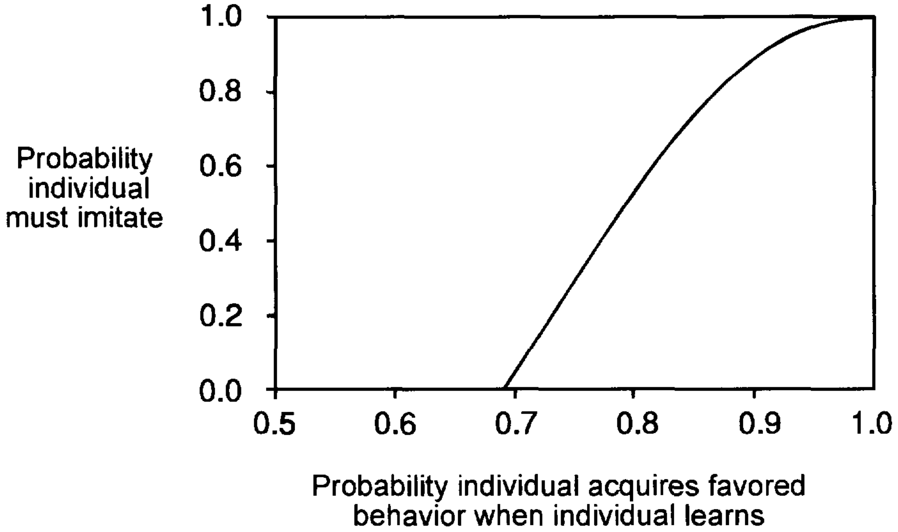

FIGURE 3. The trade-off between imitation and learning, assuming that the outcomes of learning trials are normal, distributed, with mean equal to 0.5 and variance equal to 1.0.

payoffs is a normal random variable. For one set of parameters ($\mu = 0.5$, $\sigma = 1$), the relationship between imitation and the accuracy of learning has the form shown in Figure 3. If the individual adopts a behavior any time that it yields a higher payoff during the learning trial, it will acquire the wrong behavior around 30% of the time. If it requires a larger difference in payoffs, then it can reduce the chance of such errors, but sometimes it will have to imitate. If it is sufficiently picky, it will almost never err, but it will also almost always acquire its behavior by imitation.

To model the evolution of social learning, we assume that an individual's position on this continuum is a genetically heritable trait. Suppose that most individuals use a learning rule that causes them to imitate *x*% of the time—we call these "common-type individuals." There are also a few rare "mutant" individuals who imitate slightly more often. Compared to the common type, mutants are less likely to make learning errors. Thus, when mutants learn, they have higher fitness than the common-type individuals when they learn. When mutants imitate, they have the same fitness as the common type. However, mutants must imitate more often, and imitators always have lower fitness than learners. To see why, think of each imitator as being connected to a learner by a chain of imitation. If the learner at the end of the chain learned in the current environment, then the imitator has the same chance of acquiring the favored behavior as does a learner. If the learner at the end of the chain learned in a different environment, the imitator will have a lower chance of acquiring the best behavior. Thus, the mutant type will have higher fitness if the advantage of making fewer learning errors is sufficient to offset the disadvantage of imitating more.

This evolutionary trade-off depends on how much the common type imi-

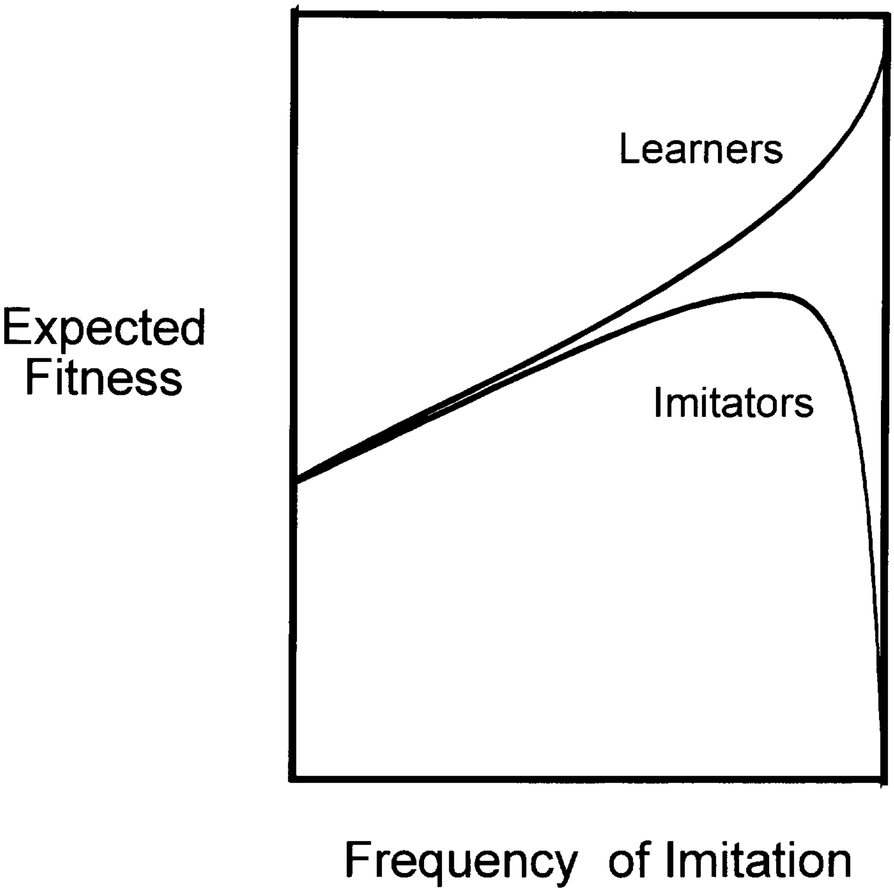

FIGURE 4. Individuals either learn or imitate according to the outcome of their learning trial. As individuals become more selective, the frequency of imitating individuals increases. This figure plots the expected fitness of individuals who imitate and those who learn as a function of the frequency of imitating individuals, assuming the outcome of learning experiments is normally distributed with mean 0.5 and variance 1.

tates. When the common type rarely imitates, the fitnesses of individuals who imitate and individuals who learn will be similar because most imitators will imitate somebody who learned, and, therefore, the fact that mutants make fewer learning errors will allow them to invade. However, as the amount of imitation increases, the fitness of imitating individuals relative to those who learn declines because increased imitation lengthens the chain connecting each imitator to a learner. Eventually an equilibrium is reached at which the common type can resist invasion by mutants that change the rate of imitation. We refer to the fraction of time that the common type imitates at equilibrium as the "evolutionary equilibrium amount of imitation."

The average fitness of a population at the evolutionary equilibrium is greater than the average fitness of individuals who do not imitate as long as the probability that the environment changes is less than one half (see Appendix 3 for a formal proof). You can get an intuitive feel for why by considering Figure 4, which plots the average fitness of imitating and learning individuals as a function of the fraction of common-type individuals who imitate. The fitness of learning individuals increases as the amount of imitation increases because learners make fewer errors. The fitness of imitating individuals also increases at first because they are imitating learners who make fewer errors. If imitation is common enough, fitness eventually declines because the population fails to track the changing environment. The first effect is apparently sufficient to lead to a net increase in average fitness at evolutionary equilibrium.

It is important to understand that this increase in average fitness is only a side effect of selection at the individual level. The evolutionary equilibrium amount of imitation does not maximize the average fitness of the population. Selection at the individual level favors more imitation than is optimal for the population because it ignores the effect on the population as a whole of increased imitation, and after a certain point this effect is deleterious.

Imitation Allows Cumulative Improvement

Imitation may increase the average fitness of learners by allowing learned improvements to accumulate from one generation to the next. So far we have only considered two alternative behaviors. Thus, learning is an either/or proposition. Many kinds of behaviors admit successive improvements towards some optimum. Individuals start with some initial "guess" about the best behavior and then invest time and effort at improving their performance. For a given amount of time and effort, the better an individual's initial guess, the better on average its final performance. Now, imagine that the environment varies, and so that different behaviors are optimal in different environments. Organisms who cannot imitate must start with whatever initial guess is provided by their genotype. They can then learn and improve their behavior. However, when they die, these improvements die with them, and their offspring must begin again at the genetically given initial guess. In contrast, an imitator can acquire its parents' behavior after their behavior has been improved by learning. Therefore, it will start its search closer to optimal behavior, and for a given amount of searching, it will achieve a better adult phenotype. Thus, if the learning cost per unit improvement is smaller for small improvements than for big ones, imitation makes learning more efficient, and therefore increases the average fitness of the population.

The following simple model illustrates this idea (a more realistic model with the same properties is analyzed in Boyd and Richerson 1985, Ch. 4). Consider an organism that lives in an environment that can be in continuum of states. For example, suppose that the population density of prey species varies. In each generation there is a chance that the environment switches to a new state (more or less prey), but also some chance that it remains unchanged. There is also a continuum of behaviors, such as the amount of effort devoted to foraging versus hunting. We measure the environmental state in terms of the optimal behavior in that environment, and assume that an individual's fitness decreases as the difference between the environmental state and its behavior value increases.

All individuals modify their behavior by learning. Each individual begins with an initial guess about the state of the environment and then experimentally modifies this behavior. In doing so, individuals reduce the difference between their behavior and the optimum behavior in the current environment. Learning is costly—individuals who devote more time and effort to experimenting suffer greater learning costs but move closer to current optimum. There are two genotypes. Learners use a fixed, genetically inherited norm of reaction as their initial guess about the environment, and they always acquire the optimum behavior. Imitators acquire their initial guess by imitating the behavior of a randomly chosen member of the previous generation. They invest much less in learning than do learners and as a result improve on their initial behavior only a small amount. However, as long as the environment does not change, the population of imitators will converge slowly toward the optimum as each generation moves toward the optimum. Thus, imitators may start their learning nearer to the optimum than do learners.

Imitators have higher fitness at evolutionary equilibrium in this model as long as (1) the environment does not change too often compared to the rate at which the population of imitators converges toward the optimum, and (2) learners suffer substantially greater learning costs than do imitators. If the environment changes slowly enough, the gradual cumulative improvement achieved by imitators will be sufficient to insure that their behavior is near the current optimum most of the time. Of course, imitators will never track the environment as accurately as learners, but if the small improvements realized by imitators are cheaper than the large improvements of learners, imitators will have higher average fitness. Because only imitators are present at such an equilibrium, imitation increases average fitness.

DISCUSSION

Culture increases average fitness if it makes the learning processes that generate new knowledge less costly or more accurate. Culture may do this in at least two ways: First, social learning allows individual learning to be selective. Individuals can learn opportunistically when it is likely to be more accurate or less costly, and imitate when conditions are less favorable. Second, social learning allows learned improvements to accumulate from one generation to the next. When learning in small steps is less costly per unit improvement in fitness than learning in large steps, the cumulative learning over many generations can increase average fitness.

These results help us understand the importance of the evolution of true imitation. There are a number of examples of social traditions in other animals. For example, some populations of chimpanzees in West Africa regularly use stone tools to crack open tough nuts, while other nearby populations never use stones to crack nuts. The stones and nuts are available to both populations, and the environments are otherwise very similar (Boesch et al. 1994). Students of social learning in non-human animals (e.g., Galef 1988; Visalberghi and Fragaszy 1990) distinguish two classes of processes that could maintain such cultural differences between different populations: *Social enhancement* occurs when the activity of older animals increases the chance that younger animals will learn the behavior on their own. Young individuals do not acquire the behavior by observing older individuals. Social facilitation could cause tool use to persist in some populations but not others, as in the following scenario: In populations in which chimpanzees use tools to crack nuts, young chimpanzees spend a lot of time in proximity to both nuts and hammer stones. Nuts are a greatly desired food, and young chimpanzees find eating nutmeats highly reinforcing. Young chimpanzees experiment with the hammers and anvils until they master the skill of opening the nuts. In populations in which chimpanzees do not use stones to open nuts, young chimpanzees never spend enough time in proximity to both nuts and hammer stones to acquire the skill. *Imitation* occurs when younger animals observe the behavior of older animals and learn how to perform the behavior by watching them. In this case, the tradition is preserved because young chimpanzees actually imitate the behavior of older chimpanzees.

Students of animal social learning have distinguished between social enhancement and imitation because the necessary psychological mechanisms are quite different. Our results suggest that this distinction is also of evolutionary importance because selective social learning and cumulative culture change are only possible when there is imitation. Social enhancement can only preserve variation in behavior that organisms can learn on their own, albeit in favorable circumstances, but it does not allow individuals to avoid learning when information is poor or costly. Even more important, only imitation allows cumulative cultural change. Suppose that on her own in especially favorable circumstances an early hominid learned to strike rocks together to make useful flakes. Her companions, who spent time near her, would be exposed to the same kinds of conditions, and some of them might learn to makes flakes too, entirely on their own. This behavior could be preserved by social enhancement because groups in which tools were used would spend more time in proximity to the appropriate stones. However, that would be as far as it would go. Even if an especially talented individual found a way to improve the flakes, this innovation would not spread to other members of the group because each individual learns the behavior anew. With imitation, on the other hand, innovations can persist as long as younger individuals are able to acquire the modified behavior by observational learning. As a result, imitation can lead to the cumulative evolution of behaviors that no single individual could invent on its own.

Recent reviews (Galef 1992; Tomasello 1990; Visalberghi and Fragaszy 1990) suggest that all known cases of animal social traditions can be explained as the result of social enhancement. If this is correct, our results explain why animal cultures seem to play such a small role in the lives of such species. It also suggests that understanding the evolution of the psychological mechanisms that allow imitation is of key importance for understanding human evolution.

We thank Alan Rogers for many useful discussions of these ideas and his careful reading of this manuscript. RB also thanks Dorothy Cheney and Robert Seyfarth for providing electricity and other facilities at Baboon Camp, where the first draft of this paper was written.

REFERENCES

Boesch, C., Marchesi, P., Marchesi, N., Fruth, B., and Joulian, F. Is nut cracking in wild chimpanzees a cultural behavior? *Journal of Human Evolution* 26:325-328, 1994.

Boyd, R., and Richerson, P.J. *Culture and the Evolutionary Process*, Chicago: University of Chicago Press, 1985.

Boyd, R., and Richerson, P.J. An evolutionary model of social learning: The effects of spatial and temporal variation. In *Social Learning: Psychological and Biological Approaches*, T. Zentall and B.G. Galef (Eds.). Hillsdale, NJ: Lawrence Erlbaum Assoc., 1988, pp. 29-48.

Boyd, R., and Richerson, P. J. Social learning as an adaptation. *Lectures on Mathematics in the Life Sciences* 20:1-26, 1989.

Galef, B.G. Imitation in animals: History, definitions, and interpretations from the psychological laboratory. In *Social Learning: Psychological and Biological Approaches*, T. Zentall and B.G. Galef (Eds.). Hillsdale, NJ: Lawrence Erlbaum Assoc., 1988, pp. 1-28.

Galef, B.G. The question of animal culture. *Human Nature* 3:157-178, 1992.

Rogers, A.R. Does biology constrain culture? *American Anthropologist* 90:819-831, 1989.

Tomasello, M. Cultural transmission in the tool use and communicatory signaling of chimpanzees. In *"Language" and Intelligence in Monkeys and Apes. Comparative and Developmental Perspectives*, S.T. Parker and K.R. Gibson (Eds.). Cambridge: Cambridge University Press, 1990, pp. 273-311.

Visalberghi, E., and Fragaşzy, D.M. Do monkeys ape? In *"Language" and Intelligence in Monkeys and Apes. Comparative and Developmental Perspectives*, S.T. Parker and K.R. Gibson (Eds.). Cambridge: Cambridge University Press, 1990, pp. 247-273.

APPENDIX 1:

Spatially Varying Environment, More Than Two Variants, Learning Errors

Consider an organism that lives in a spatially varying environment in which there are a large number of islands. A different behavior is favored on each island so that the fitness of behavior i on island j is

$$W_i = \begin{cases} W_0 + D \text{ in environment } i \\ W_0 - D \text{ in environment } j \end{cases} \qquad \text{A1.1}$$

There are two genotypes:

Learners = Discover locally optimal behavior with probability $1 - e$.
Imitators = Imitate a randomly chosen individual from the previous generation.

After learning and imitating, a fraction m of the individuals on each island emigrate and are replaced by individuals drawn from all other islands at random. Because the number of behaviors is large, the frequency of the favored behavior among immigrants is approximately zero.

After migration, selection occurs. We assume that selection is weak so that the frequency of innovators and imitators is the same on all islands. Then let

q = frequency of imitators on the focal island.
p = frequency of the locally favored behavior among imitators.

The probability that an imitator encounters a single individual who has the locally optimal behavior is $(1 - q)(1 - e) + qp$, and thus the frequency of the locally optimal trait among imitators after imitation, p', is

$$p' = (1 - q)(1 - e) + qp \tag{A1.2}$$

And after migration the frequency of the favored behavior among imitators, p'', is

$$p'' = (1 - m)[(1 - q)(1 - e) + pq] \tag{A1.3}$$

Thus, there is a unique stable equilibrium frequency of the locally favored variant, $\tilde{p}$

$$\tilde{p} = \frac{(1 - m)(1 - e)(1 - q)}{1 - q(1 - m)} \tag{A1.4}$$

The average fitness of learners is $W_L = W_0 + D(1 - 2e) - C$ where C is the cost of individual learning. The average fitness of imitators, $W_I = W_0 + D(2\tilde{p} - 1)$. If imitators are rare ($q \approx 0$), then the equilibrium frequency of the favored variant among rare imitators is approximately $(1 - e)(1 - m)$, the same frequency as among learners, and since imitators incur no learning cost, they increase in frequency. If imitators are common ($q \approx 1$), then the equilibrium frequency of the favored variant is zero, and therefore imitators have lower fitness than do learners as long as learning pays $[D(1 - 2e) > C]$. Since $\tilde{p}$ is a monotonically decreasing function of q, there is a unique stable equilibrium value of q at which imitators have the same fitness as learners.

APPENDIX 2:

Imitators Can Identify Learners

Consider an organism that lives in an environment that can be in one of two states. Each generation there is a probability γ that the environment switches from one state to the other. There are two behaviors with fitnesses as given in the following table:

	Environment 1	Environment 2
Behavior 1	$W_0 + D$	$W_0 - D$
Behavior 2	$W_0 - D$	$W_0 + D$

There are two genotypes:

Learners = Always acquired the best behavior in the current environment but at a cost C.

Imitators = Observe n individuals after learning. If there is a learner among these individuals, imitators acquire the best behavior in the current environment. Otherwise they copy a random individual from within the group.

And let q equal the frequency of imitators, and p the frequency of the currently favored behavior among imitators. Assume that selection is sufficiently weak so that the effect of selection on cultural evolution can be ignored (i.e., on dynamics of p), and genetic evolution (the dynamics of q) responds to the stationary distribution of p.

Then the frequency of the currently favored behavior after learning and imitation is

$$p' = \begin{cases} 1 - q^n + q^n p & \text{if no environmental change} \\ 1 - q^n + q^n(1 - p) & \text{if environment changes} \end{cases} \qquad \text{A2.1}$$

Suppose at some time t the probability density for p is $f_t(p)$ with mean P_t. Then the mean of $f_{t+1}(p)$ given by

$$P_{t+1} = \int [(1 - \gamma)(1 - q^n + q^n p) + \gamma(1 - q^n + q^n(1 - p))] f_t(p) dp \qquad \text{A2.2}$$

where γ is the probability that the environment switches states. Integrating and simplifying yields the following recursion for P_t

$$P_{t+1} = 1 - q^n + q^n[(1 - 2\gamma)P_t + \gamma] \qquad \text{A2.3}$$

Thus the equilibrium value of mean frequency of the favored behavior is:

$$P = \frac{1 - q^n + q^n\gamma}{1 - q^n(1 - 2\gamma)} \qquad \text{A2.4}$$

The average fitness of learners is $W_L = W_0 + D - C$, which is independent of changes in the environment. The average fitness of imitators once P_t has reached its equilibrium value is $W_I = W_0 - D(2P - 1)$. The frequency of imitators will increase whenever $W_I > W_L$. Substituting the expression for P given in equation A2.4 and solving for q yields the following inequality:

$$q < q^* = \left(\frac{C/D}{2\gamma(1 - C/D) + C/D} \right)^{1/n} \qquad \text{A2.5}$$

Thus q^* is a unique stable equilibrium value for the frequency of imitators, and at this frequency the average fitness of imitators and learners is equal.

APPENDIX 3:

Selective Learning

Consider an organism that lives in an environment that can be in one of two states. Each generation there is a probability γ that the environment switches from one state to the other. There are two behaviors with fitnesses as given in the following table:

	Environment 1	Environment 2
Behavior 1	$W_0 + D$	$W_0 - D$
Behavior 2	$W_0 - D$	$W_0 + D$

Each individual performs a learning trial in which it estimates the payoff of each behavior in the current environment. The difference between the payoff of the currently favored behavior and that of the alternative behavior observed by each individual is an independent, normally distributed, random variable, x, with mean equal to m, and variance equal to 1. The mean, m, is positive because, on average, the currently favored behavior yields a higher payoff in the current environment. All individuals use the learning rule:

Outcome of Learning Trial	Decision
$x > d$	adopt favored behavior
$d > x > -d$	imitate
$-d > x$	adopt other behavior

The threshold parameter d determines how selectively individuals learn. Individuals regard trials that yield positive outcomes greater than d as decisive evidence that the environment is in the state that is currently favored, and trials in which x is less than $-d$ as decisive evidence that the environment is in the other state. When a trial produces an outcome in between d and $-d$, it is indecisive and individuals imitate.

The value of d is a genetically heritable trait. At any time there are two genotypes present in the population. Most of the population has $d = d^*$, but there are a very few rare mutants who have $d = d^* + \delta d$. We seek to determine the values of d^* that can resist invasion by mutants with slightly smaller or slightly larger values of d. Such continuous ESS solutions often yield the same outcome as genetically more realistic models.

Let v be the frequency of the favored behavior in the population. Assume that selection is sufficiently weak so that the effect of selection on cultural evolution can be ignored (i.e., on dynamics of v) and genetic evolution responds to the stationary distribution of v. Finally, let $p_1(d) = \Pr(x > d)$, $p_2(d) = \Pr(x < -d)$, and $L(d) = 1 - p_1(d) - p_2(d)$; $p_1(d)$ is the probability of correctly choosing the currently favored behavior, $p_2(d)$ is the probability of mistakenly choosing the other behavior, and $L(d)$ is the probability of imitating. Then the frequency of the favored variant in the next generation, v', is:

$$v' = \begin{cases} vL(d^*)+p_1(d^*) & \text{if no change in environment} \\ (1-v)L(d^*)+p_2(d^*) & \text{if environment changes} \end{cases} \qquad \text{A3.1}$$

Suppose at some time t the probability density for v is $f_t(v)$ with mean V_t. Then the mean of $f_{t+1}(v)$ given by

$$V_{t+1} = \int[(1 - \gamma)(vL + p_1) + \gamma((1 - v)L + p_2)]f_t(v)dv \qquad \text{A3.2}$$

Integrating and simplifying yield the following recursion for V_t

$$V_{t+1} = (1 - 2\gamma)(V_t L + p_1) + \gamma \qquad \text{A3.3}$$

Thus the equilibrium value of the mean frequency of the favored behavior is:

$$V = \frac{(1 - 2\gamma)p_1 + \gamma}{(1 - 2\gamma)(p_1 + p_2) + 2\gamma} \qquad \text{A3.4}$$

The fitness of the common genotype averaged over the stationary distribution of v is

$$\overline{W}(d^*) = W_0 + D[VL(d^*) + p_1(d^*)] - D[(1 - V)L(d^*) + p_2(d^*)] \qquad \text{A3.5}$$

and the fitness of the mutant type is

$$\begin{aligned}\overline{W}(d^* + \delta d) = {} & W_0 + D[VL(d^* + \delta d) + p_1(d^* + \delta d)] - \\ & D[(1 - V)L(d^* + \delta d) + p_2(d^* + \delta d)]\end{aligned} \qquad \text{A3.6}$$

Thus because δd is small, the difference in fitness between the mutant and common types, δW, is

$$\delta W = D\left((2V - 1)\left(\frac{\partial L}{\partial d}\right)_{d^*} \delta d + \left(\frac{\partial p_1}{\partial d}\right)_{d^*} \delta d - \left(\frac{\partial p_2}{\partial d}\right)_{d^*} \delta d\right) \qquad \text{A3.7}$$

Setting $\delta W = 0$, substituting the expression for V given in A3.4, and simplifying yields the following necessary condition for the ESS,

$$0 = (1 - 2\gamma)\left(\left(\frac{\partial p_1}{\partial d}\right)_{d^*} p_2(d^*) - \left(\frac{\partial p_2}{\partial d}\right)_{d^*} p_1(d^*)\right) + \gamma\left(\left(\frac{\partial p_1}{\partial d}\right)_{d^*} - \left(\frac{\partial p_2}{\partial d}\right)_{d^*}\right) \qquad \text{A3.8}$$

Given that x is normal with a known mean and variance, this equation can be solved numerically for the value of d^*.

We now prove that the average fitness of a population at the ESS value of d, d^*, is greater than the average fitness of a population with no imitation (i.e., $d = 0$) whenever $m > 0$ and $\gamma < 1/2$. It follows from A3.8 that when $\gamma = 1/2$ then $d^* = 0$ and, therefore, that $\overline{W}(d^*) - \overline{W}(0) = 0$. Next, we show that $\overline{W}(d^*) - \overline{W}(0)$ is a monotonically decreasing function of γ as long as m is positive. Compute

$$\begin{aligned}\frac{\partial}{\partial \gamma}(\overline{W}(d^*) - \overline{W}(0)) = {} & 2\frac{\partial V}{\partial \gamma}L(d^*) + \frac{\partial d^*}{\partial \gamma}\left\{(2V - 1)\left(\frac{\partial L}{\partial d}\right)_{d^*} + \right. \\ & \left. \left(\frac{\partial p_1}{\partial d}\right)_{d^*} - \left(\frac{\partial p_2}{\partial d}\right)_{d^*}\right\}\end{aligned} \qquad \text{A3.9}$$

But the ESS condition (A3.7) guarantees that the term in braces on the right hand side of A3.9 is zero. Thus

$$\frac{\partial}{\partial\gamma}(\overline{W}(d^*) - \overline{W}(0)) \propto \frac{\partial V}{\partial\gamma} = p_2(d^*) - p_1(d^*) +$$

$$(1 - 2\gamma)\frac{\partial d^*}{\partial\gamma}\left\{\gamma\left(\frac{\partial p_1}{\partial d} - \frac{\partial p_2}{\partial d}\right)_{d^*} + (1 - 2\gamma)\left(\frac{\partial p_1}{\partial d}p_2 - \frac{\partial p_2}{\partial d}p_1\right)_{d^*}\right\} \qquad \text{A3.10}$$

Once again the ESS condition guarantees that the term in braces on the right hand side of A3.10 is zero, and since $p_1(d^*) > p_2(d^*)$ for $m > 0$, it follows that the average fitness of an ESS population is greater than the fitness of a population with no imitation as long as $\gamma < 1/2$.

APPENDIX 4:

Cumulative Learning

Consider an organism that lives in an environment that can be in a continuum of states. Each generation there is a probability γ that the environment switches from its current state to a new state drawn at random from a probability distribution with mean equal to zero and variance equal to H. There is a probability $1 - \gamma$ that the environment will remain unchanged. There is also a continuum of behaviors. In each environment, fitness is a gaussian function of behavior so that there is a unique optimum behavior θ_t. We choose to measure the state of the environment as the optimal behavior in that environment. All individuals modify their behavior by learning so that the difference between their behavior and optimum behavior in the current environment is reduced. There are two genotypes:

Learners = Acquire the optimal behavior. Learning costs reduce fitness by a factor e^{-C_L}.

Imitators = Imitate a randomly chosen individual from the previous generation, and then adjust their behavior a small fraction, a ($a << 1$) by learning. Learning costs reduce fitness by a factor e^{-C_I}.

Suppose most individuals in a population are imitators, but that there are a small number of rare learners. Because they always acquire the optimal behavior, the expected fitness of learners is simply:

$$W_L = \exp(-C_L) \qquad \text{A4.1}$$

and the expected fitness of copiers is:

$$W_I = \exp[-(1 - a)^2(Z_t - \theta_t)^2 - C_I] \qquad \text{A4.2}$$

where Z_t is the behavior of imitators during period t, which will change from period to period according to the following recursion.

$$Z_{t+1} = a\theta_t + (1 - a)Z_t \tag{A4.3}$$

Thus the behavior of imitators will converge toward the current optimum at a rate a. When the environment changes, it will converge toward a different value. Assume that selection is weak enough that changes in gene frequency respond to the stationary distribution of Z_t. Thus imitation is evolutionarily stable if

$$-(1 - a)^2E\{(Z_t - \theta_t)^2\} - C_I > -C_L \tag{A4.4}$$

where the expectation is taken with respect to the joint stationary distribution of θ_t and Z_t.

$$E\{(Z_t - \theta_t)^2\} = E\{Z_t^2\} - 2E\{Z_t\theta_t\} + E\{\theta_t^2\} \tag{A4.5}$$

To compute $E\{Z_t\theta_t\}$ multiply both sides of A4.3 by θ_{t+1}.

$$\theta_{t+1}Z_{t+1} = a\theta_t\theta_{t+1} + (1 - a)Z_t\theta_{t+1} \tag{A4.6}$$

Taking the expectation of both sides yields

$$E\{\theta_{t+1}Z_{t+1}\} = a[(1 - \gamma)V + \gamma 0] + (1 - a)[(1 - \gamma)E\{\theta_t Z_t\} + \gamma 0] \tag{A4.7}$$

The moments of the stationary distribution are constant, and thus setting $E\{Z_{t+1}\theta_{t+1}\} = E\{Z_t\theta_t\}$ and solving yields:

$$E\{Z_t\theta_t\} = \frac{a(1 - \gamma)V}{1 - (1 - a)(1 - \gamma)} \tag{A4.8}$$

To compute $E\{Z_t^2\}$ square both sides of A4.3.

$$Z_{t+1}^2 = a^2\theta_t^2 + 2a(1 - a)Z_t\theta_t + (1 - a)^2Z_t^2 \tag{A4.9}$$

Again taking the expectation of both sides, setting $E\{Z_{t+1}^2\} = E\{Z_t^2\}$, and substituting the expression for $E\{Z_t\theta_t\}$ yields:

$$E\{Z_t^2\} = \frac{a[1 + (1 - a)(1 - \gamma)]}{(2 - a)[1 - (1 - a)(1 - \gamma)} \tag{A4.10}$$

Substituting the expressions for $E\{Z_t\theta_t\}$ and $E\{Z_t^2\}$ into A4.5 and simplifying yields:

$$E\{(Z_t - \theta_t)^2\} = \frac{2\gamma V}{(2 - a)[1 - (1 - a)(1 - \gamma)]} \tag{A4.11}$$

Substituting this expression into A4.4, ignoring terms of order a^2, and simplifying yields the following condition for imitation to be an ESS.

$$\left(\frac{\delta}{1-\delta}\right)a > \gamma \tag{A4.12}$$

where $\delta = \frac{C_L - C_I}{V}$ is the fitness advantage of imitators due to lower cost learning measured in units of V the average log fitness increase of learners due to learning. Because learning would not be favored by selection for learners if $V < C_L$, we know that $\delta < 1$. Recall that a is the rate at which imitators converge toward the current optimum. Thus, the ESS condition, A4.12, says that the rate of environmental change must be less than the rate at which imitators converge toward the current optimum as modified by the term in parentheses. This term is greater than one when the learning cost advantage of imitators is a large fraction of the total benefit of learning, and less than one when the learning cost advantage of imitators is relatively small.