RICHARD MCELREATH

Social Learning and the Maintenance of Cultural Variation: An Evolutionary Model and Data from East Africa

ABSTRACT Human societies maintain between-group variation despite mixing of people and ideas. In order for variation to remain, migrants or their children must preferentially adopt local norms, customs, and beliefs. Yet the details of how cultural variation is maintained, despite mixing, remain unknown. This article addresses this problem by using a simple model of the evolution of cultural learning to interpret the results of a study of cultural variation in a small region of East Africa. I argue that the manner in which migrants of two diverse regions adapt to local beliefs and behavior depends on the costs and accuracy of learning in each domain. Observational studies are never definitive tests of any hypothesis, but these results suggest that conclusions about the significance of cultural learning for understanding individual attitudes and behavior depend strongly upon the domain of investigation. [Keywords: cultural evolution, social learning, East Africa, cultural variation]

STUDY OF THE DIVERSITY of human adaptation, social organization, and belief has occupied more than a century of anthropological description and explanation. Most social scientists are convinced that this diversity arises from a number of learning strategies, both simple and complex, and that sophisticated social learning in particular plays a key role in transmitting variation in behavior between generations. Many social scientists refer to behavior that is at least partly acquired via social learning as culture (Kroeber and Kluckhohn 1952). Under this definition, few deny the existence of important cultural differences among individuals and societies.

How particular differences form and persist, however, is a matter of more debate. As individuals migrate between cultural regions or interact with their culturally different neighbors, unbiased imitation and socialization will lead to the blending and the eventual disappearance of cultural variation (McElreath et al. 2003). In order for variation to be maintained, migrant individuals or their children must either (1) learn local behaviors and beliefs on their own or (2) be more likely to adopt the behaviors and beliefs of nonmigrants.

First, some social scientists ask us to consider that large portions of "culture" are actually not culture at all but, instead, variation in behavior evoked by individual circumstances (Tooby and Cosmides 1992). Under this view, many apparent cultural differences form and persist because of differences in the physical and social environments individuals experience. Migrants (or their children) experience a new economy, ecology and social world, and figure out how to behave. Social learning may still play a role in transmitting solutions to successive generations, but forms of individual learning "evoke" (Tooby and Cosmides 1992) new behavior when circumstances change, because of individuals migrating or to changes in the local context itself. When such mechanisms are important, traditions and behavioral differences in general persist only because differences in relevant features of the environment persist. Much of the classic work in cultural ecology adopts this perspective for some proportion of cultural variation, either implicitly or explicitly (Goldschmidt 1976). Economists, the most influential social scientists by far, also generally rely on environmental and external market explanations of behavioral differences (Becker 1976).

A second possibility is that the acquisition of behavior via social learning itself leads to the formation and persistence of variation among groups—from small entities like families to very large ones like multilingual states—perhaps largely independent of underlying ecological and economic differences. For example, if migrants and their children preferentially imitate the most common beliefs and behaviors, then cultural variation can be stable despite substantial mixing of people and ideas (Boyd and Richerson 1985). Theories of this general kind are very common and often poorly

American Anthropologist, Vol. 106, Issue 2, pp. 308–321, ISSN 0002-7294, online ISSN 1548-1433.

specified. Anthropology as a field has done a poor job constructing explicit models of cultural evolution, despite the fact that culture and human variation are among its central concepts. A small group of anthropologists, however, have taken on the task of constructing explicit models of cultural evolution (Boyd and Richerson 1985; Cavalli-Sforza and Feldman 1981; Durham 1991; Odling-Smee et al. 2003). Joseph Henrich and Richard McElreath (2003) provide a recent review. These models tackle how social learning can both build complex behavioral and technological adaptations and produce cultural inertia, variation in beliefs, and maladaptive behaviors. Under these theories, differences may form and persist (or not) for a variety of reasons, but in each case appreciation of how, and from whom, people acquire their culture is key. Evolved psychology and individual rationality is important here in understanding how individuals strategically acquire their culture, but less so in preparing individuals to learn any specific behavior.

Most would agree that both of the perspectives above are sometimes correct (Laland and Brown 2002): Some observed cultural variation is the result of individuals discovering what works in local circumstances, while some is the product of patterns of imitation and other forms of cultural learning. The questions are: How often and when? In this article, I approach the problem of how cultural differences persist by first outlining a general theory of the relative benefits of individual and social learning. Then I present relevant results from an ethnographic study of the maintenance of behavioral variation, patterned after the classic studies of R. B. Edgerton (1971) and W. R. Goldschmidt (1976) on how ecology relates to cultural variation. Although their work demonstrated general correlations between ecological and cultural differences, it failed to exploit the variation within communities to help understand how variation is maintained, despite the mixing of peoples from different regions. Using this natural variation in my own field site, I argue that the general models of learning make useful predictions that are consistent with the observations from the field. These results are important for understanding how differences in human behavior form and persist, because they help us understand how individuals strategically acquire their behavior and beliefs, which in turn allows prediction of the more long-term patterns of cultural diversity, including which changes in the physical and social environments are likely to enhance or erode variation.

THE DESIGN OF LEARNING STRATEGIES

In this section, I outline a quantitative model of the evolution of individual and cultural learning. An evolutionary model is necessary, because the sophisticated learning abilities which make cultural variation possible are the results of a long history of selection acting on genes. This does not mean that cultural variation itself is a product of genetic variation, nor does it suggest that culture must ultimately serve genetic interests—as even the simplest models of the evolution of cultural learning show (see Boyd and Richerson 1985; Rogers 1988). This model is instead an exercise in deducing the consequences of assuming that individuals can acquire their behavior in different ways and asking how people might employ these different learning strategies under different circumstances. The goal is to derive qualitative predictions that we can apply to our observations. Setting out the theory in this way also allows us to see exactly what is being assumed, something that is usually not easy with purely verbal models. I use the model to generate predictions about the conditions under which different learning mechanisms are likely to influence how individuals acquire behavior. These predictions in turn suggest what variables will influence patterns of behavioral variation and how cultural variation might be maintained in any given setting.

Psychologists and anthropologists interested in the evolution of learning distinguish between *individual* and *social* learning (see chapters in Heyes and Galef 1996). Individual learning occurs when people (or other animals) interact with their environment and adjust their behavior in light of new information. Social learning occurs when people are more likely to learn a behavior because of the behavior of other individuals. This includes both complex forms of social learning, like imitation, and simple forms in which other individuals alter the environment such that individuals practicing forms of individual learning are more likely to figure something out. In what follows, I use "social learning" to refer to forms that require psychological mechanisms distinct from individual learning.

A Simple Model of Individual and Social Learning

Here I develop a simple model of the evolution of social learning that explicitly allows for mixed learning strategies as well as the costs of and errors in learning. My purpose is to see how variation in the costs and accuracy of learning affects which learning strategies will arise. Henrich and Boyd (1998) simulate a quite general model that includes accuracy and produces the same qualitative conclusions. The model here should be thought of as a simplified argument designed to include the bare minimum necessary to understand how this family of models generates the intuitions I will draw from them. I outline the assumptions and conclusions here, and I relegate the mathematical details to the appendix.

Imagine a large number of individuals who all face the same decision, for which there is a very large number of different possible behaviors. This may represent a decision such as what mix of crops to plant, how to rotate them, and how to divide labor between crops and livestock. The optimal behavior is determined by many exogenous variables and, therefore, may be difficult to determine. Let b be the average value of behaving optimally under current environmental conditions. Behaving nonoptimally entails a payoff of zero. Thus b actually measures the extent to which behaving optimally is better than behaving otherwise. Assume that individuals can, through interacting with the environment, sometimes individually acquire the optimal

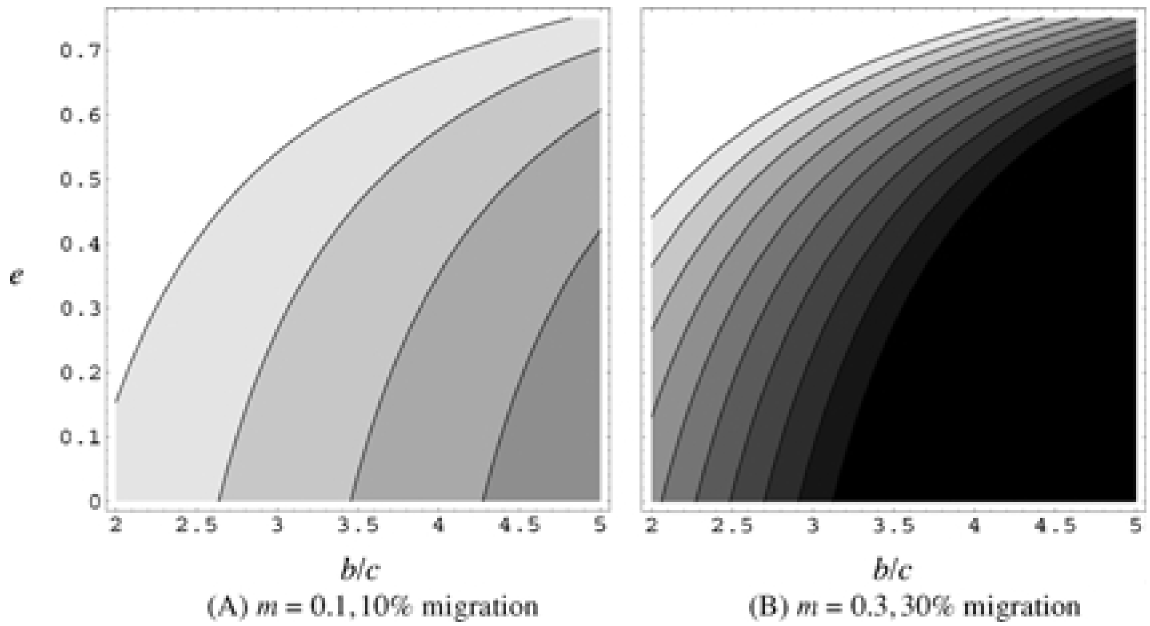

FIGURE 1. Evolutionary stable proportion of individual learning, for two probabilities of emigration (m), as functions of the rate of learning error (e) and the ratio of the benefits of behaving optimally to the costs of learning (b/c). At the top of each plot, individual learning is very error-prone, because information is unreliable or the problem is too difficult for individuals to solve. At the left side of each plot, individual learning is expensive in terms of effort or mistakes. Black indicates strong reliance on individual learning. White indicates strong reliance on social learning. Each contour is a ten percent change in the proportion of social learning at equilibrium. (A) $m = 0.1$. A substantial reliance on social learning is favored over most of the parameter space, even when costs of learning (c/b) are low and accuracy ($1 - e$) is high. (B) $m = 0.3$, very high rate of migration. Individual learning excludes pure social learning for a good range of values, when either costs are low or accuracy is high. When costs are high (b/c low) or accuracy is low (e high), however, social learning may still dominate.

behavior, but e of the time they cannot and are stuck with what people have done before—whether it is optimal or not. Individual learning is costly, however, because of opportunity costs and potentially dangerous events while exploring the environment, resulting in an average cost c.

Assume now there are a number of environments, and the population is split among them. Individuals emigrate each generation with probability m. This captures the nearly universal phenomenon of mixing of peoples due to intermarriage, capture, and exchange of ideas. Individuals, both natives and those new to an environment, must decide each generation how to acquire their behavior. In the appendix, I show how the above assumptions lead to an expression for the mix of individual and social learning that will result from natural selection on genes influencing this mix. This expression allows us to vary the costs and accuracy of learning and generate qualitative predictions about how much individuals might rely upon individual or social learning under different circumstances. Over much of the parameter space, a stable mix of individual and social learning exists (Figure 1). Social learning is profitable only if there is not too much migration (m is not too large). Provided this is true, social learning does better as the costs of individual learning (c) and the rate of learning errors (e) increase. If individual learning is sufficiently cheap and accurate, there is no problem for social learning to solve, and individual learning will either be very common or exclusively present. Some pure social learning can evolve under a wide range of conditions because pure social learning exploits already paid costs of individual learning needed to generate novel adaptations.

Model Summary and Predictions

The model above suggests that the amount of pure social learning, as a proportion of total learning, will increase as the costs of individual learning increase and the accuracy of individual learning decreases. This qualitative result is unlikely to depend on the precise assumptions of the model, because the tradeoff at work is a very basic one, common to a wide range of models (Boyd and Richerson 1985; Rogers 1988). One key advantage of social learning lies in avoiding the costs of individual trial and error. The price is that pure social learners need to be at least partly credulous, and so risk adopting outdated or inaccurate behavior. These benefits and costs balance at a proportion of pure social learning, which is often rather high. However, when learning on one's own is very inexpensive and highly accurate, social learning may be quite rare. When learning is costly and error prone, social learning should be much more frequent.

It is easy to interpret the costs and accuracy as taking different values, or ranges of values, for different tasks and problems. Instead of being prepared to employ individual learning with the same likelihood in all instances, people are probably sensitive to the costs and accuracies in different contexts. For example, tasting mushrooms has potentially very high costs, because of poisoning, but the accuracy might be quite high, as a nonpoisonous mushroom is easily identified once sampled. In this case, the costs are high, as

TABLE 1. Example domains and qualitative differences in costs of individual learning (time spent, price of experimentation) and accuracy of learning (chance of process leading to better/best solution). The models in the main text suggest social learning will be more important (adaptive and frequent) when costs are high (cells 2 and 4) or accuracy is low (cells 3 and 4). Only in cell 1 is individual learning expected to be dominant. Social learning is expected to especially important in cell 4, in which both accuracy is low and costs are high.

	Costs of individual learning	
Accuracy	Low	High
High	1. **Choosing tool raw materials:** Materials often easily acquired, abundant, so experiments cheap; effectiveness easy to evaluate outside critical tasks.	2. **Identifying dangerous plants/animals:** Experiments potentially deadly; experiments very diagnostic, however.
Low	3. **Predicting the weather:** Relevant information largely free, but mostly uncorrelated with outcomes or insufficient to make an accurate judgment.	4. **Choosing a crop/mate/place to live:** Experiments time-consuming and expensive; many variables important, so outcomes highly variable.

well as the accuracy, and we would expect a good amount of pure social learning. Deciding whether to use steel or stone to make an axe, however, is likely to be cheap and rather accurate. In this case, we would expect more individual learning. Table 1 gives examples of different real-world problems which vary qualitatively in costs of learning and degree of accuracy. Of course, what individuals may be learning, either individually or socially, is how to make decisions in these domains, and not specific choices. For example, one might be taught what cues to attend to in predicting the weather and even how to integrate multiple cues.

But how do individuals estimate the costs or accuracy of learning? This is another important problem to be solved. One way in which individuals might experience accuracy is when there is a strong correlation between behavior and outcome. If only some mushrooms of a deadly variety are poisonous, then experiments are less diagnostic and the accuracy of individual learning is reduced. In general, variation in outcomes might serve as a cue to the low quality of information available via interacting with the environment. Every farmer knows that many factors influence yield, leading to high variance in yields for apparently similar crops and methods, and so distilling the signal from the noise in such information is difficult. The costs of learning might be assessed by estimating the time and energy put into experiments and judging the risks involved in exploration. For example, it is obvious to most of us that trying out a career for a few years is a very costly form of individual learning. A series of anthropologically informed experiments addressing these hypotheses would go a long way towards linking individual experience to the theory.

It goes without saying that the arguments above only hold when all other things are equal. That is, any number of factors such as ritual or symbolic importance of a behavior or technology could alter the predictions. Yet general theories are just as necessary for understanding cultural variation as is appreciating the details of each case. The consequences of the model for understanding the maintenance of cultural variation in human societies lie in recognizing that the differential use of individual and social learning will lead to different predictions about the processes that lead migrant individuals or their children to become behaviorally like locals. When individual learning is important, variables related to the payoffs and nature of the specific problem will best predict individual adaptation to local norms of behavior. As individuals interact with their new physical and social environments, they learn about behavior and may come to resemble neighbors not because of imitation but, rather, because they now experience similar reinforcements. But when social learning is important, individual payoffs and other aspects of the problem may only weakly determine behavior. Migrants may realize that their previous behaviors are poorly suited to their new environment, but since individual learning is ineffective or costly, they will learn socially from neighbors and bypass individual reinforcements, as a best strategy. If individual learning is weak enough, then traditions can persist in such domains despite important changes in the environment.

AN EMPIRICAL STUDY OF THE MAINTENANCE OF CULTURAL VARIATION

In order to investigate how cultural variation is maintained, I selected a natural experiment that allowed me to address the maintenance of intraethnic cultural variation predicated on ecological differences. Not all or even most cultural variation is necessarily patterned after the physical environment, but since at least some cultural variation promotes individual survival and reproduction in terms of the local environment, it makes sense to formulate and test hypotheses about how different cultural variants may be adaptations to different ecologies and economies. It also makes sense to use regional comparisons within closely related groups to increase the chance that observed differences are derived in each case and to ease interpretation of the causes of variation (Johnson 1991). Differences measured against a vast background of similarity admit fewer possible explanations than differences measured against a background of immense difference.

Walter Goldschmidt (1976) and Robert Edgerton's (1971) study of four East African tribes, each varied internally by ecology and economy, is the best ethnographic research to date addressing the correlations between ecological and cultural variables. Edgerton used a standard questionnaire in eight communities: four predominantly pastoralist and four predominantly farming, among four tribes—the Kamba, Pokot, Hehe, and Sebei. While ethnic group and language group predicted most of the cultural variation in the whole data set, several important aspects of personality and culture correlated strongly with ecology and economy, irrespective of ethnic and language group.

Unfortunately, their study did not record much individual-level data, and so it could not address how overall community-level differences, correlated with ecology, were generated or maintained. Edgerton could not answer, for example, if farmers living in pastoralist communities were more like farmers or pastoralists of their ethnic group, because he did not know who in the sample owned livestock and who did not, but only where each informant lived at the time of the interview. Answering questions of this kind is important, because there are many ways that variation between farmers and pastoralists could be maintained, and sorting among them focuses us on more specific hypotheses to test and effects to measure. There might be a developmental period during which some aspects of belief and behavior are fixed, meaning that variation is maintained partly because of early socialization. It may also be true that parents are important in early cultural transmission, such that the children of migrants are different from the children of natives, but that peer influence increases with age, such that adult patterns of variation are maintained despite the flow of persons between regions. Another important factor to consider is the kind of household (farming or herding) in which an individual lives or, rather, the kind of household in which most individuals in the local community live.

To begin to address these questions, I selected a field site in Tanzania patterned after the Edgerton (1971) study, with a strong ecological gradient, along which a single ethnic group is economically differentiated into two groups: (1) predominantly farmers and (2) predominantly pastoralists. Using a subset of Edgerton's instrument, I replicated some overall differences between farmers and pastoralists and then used individual-level data to shed light on the processes that have allowed these differences to persist, despite substantial contact and flow of individuals between the regions. Migrants between farming and pastoralist economies provide the crucial data to evaluate the predictions derived above. If individual learning is important in a given domain, individual reinforcements arising from the nature of the problem should predict changes in individual behavior. If social learning is dominant, then individual reinforcements will be less relevant than features of the community in predicting the behavior of migrants.

Ethnographic Setting

The Usangu Plains is a 15,000 square kilometer region of Southwest Tanzania that is home to a number of ethnic groups living throughout a gradient between wet and dry environmental zones (Hazlewood and Livingstone 1978; Pipping 1976). In the wetter regions of the south, annual rainfall ranges from 600–1000 millimeters. The dry, seasonally flooded range lands to the north vary widely in rainfall, from 200–600 millimeters annually, with much more interannual variation at specific locales.

Individuals who identify themselves as Sangu vary tremendously in economic organization, as a consequence of the ecological variation in the plains. To this day the single largest ethnic group in the Usangu Plains, the Sangu originated from a mixture of Bantu peoples present in the region in the late 1800s and early 1900s, when they united under a hereditary chief and began raiding their neighbors for livestock and taking slaves (Shorter 1972; Wright 1971). At the peak of their power, the Sangu were wealthy cattle pastoralists who wielded considerable military might. Presently, most Sangu are farmers, although probably a hundred households still keep herds in the plains.

Since 1997, I have been doing research in both the agricultural and pastoralist zones of Usangu. In the areas surrounding and including the village of Utengule, the vast majority of households exclusively farm, although some households keep moderate-sized herds. Residents live in very closely spaced settlements, where often there is less than ten meters between homes. Most farm on small plots of land, between one-half and two acres in size, outside the village. The major pastoralist regions begin about thirty kilometers north and east of Utengule. The areas surrounding and including the villages of Ukwaheri, Uyaule, Mjenje, and Upagama are less villages than collections of small patrilineal clan-based communities. Household compounds are very scattered. Distances of one to two kilometers are common, although households within a clan tend to reside closer. Some residents of the drier regions do not own cattle. Many, however, own at least some livestock, and those with larger herds (typically more than 20 cattle) practice transhumance. Every household farms at least an acre of corn, although the land in this region is much less productive than in the farming region.

Methods

Diversity of household economies within each region allows me to estimate separate influences of economy and community. Changes in household economy, as the result of intrahousehold changes and migrations, allow me also to estimate the effects of the economy individuals grew up in. I interviewed 169 Sangu individuals, 83 in the farming region and 86 in the pastoralist region, over a nine-month period, sampling across age classes, with approximately equal numbers of men and women, to capture both average regional differences and the dispositions of individuals who have changed household economies during their lives. I conducted interviews in Swahili. Aside from a local assistant and the informant, only nursing infants were allowed to be present during interviews. During interviews, I recorded demographic and economic data on each individual, as well as personal histories of residence, and economic information for father's household. For a subsample, I validated herd sizes and acres under cultivation by observation and found informant reports to be very accurate. For all

TABLE 2. Interview questions discussed in the article, as well as their coded responses.

Domain	Question	Response Codes
Preference for friends	Think of a person who has many friends [*marafiki*], but very few kin [*ndugu*], for example brothers or cousins. Now think of a person who has very few friends, but many kin. Which kind person would you rather be?	Friends, Kin
Respect for elders	Think of a household consisting of a father [*mzee*] and his sons [*vijana wake*]. Now, imagine that father has committed an error of judgment. When can one of the sons tell the father he is wrong?	Whenever, Never
Witchcraft	Imagine someone owns a number of cattle. One day, 10 of those cattle suddenly die. What is this person likely to think caused the deaths?	Other misfortune, Witchcraft

herd sizes, I also checked reports with estimates from two other unrelated individuals.

From Edgerton's (1971) results, I selected to investigate those domains that showed the greatest variation in response between farmers and pastoralists in order to maximize statistical power. I report the results from the three domains that replicated Edgerton's findings, showing strong differences between Sangu farmers and pastoralists. Therefore, my results shed light on the maintenance of behavioral variation. Table 2 summarizes the interview questions.

Domain 1: Preference for Friends or Kin

In four different ethnic groups in East Africa, Edgerton (1971) previously found that farmers, more often than pastoralists, valued friends over kin. One potential explanation for this finding is that kin are much more reliable cooperating partners when community membership is shifting, as in partly nomadic pastoralist communities. Models of the evolution of reciprocity (Axelrod and Hamilton 1981; Trivers 1971) and large-scale cooperation (Boyd and Richerson 1992) suggest that cooperation among unrelated individuals, such as friends, is stable only when pairings are sufficiently likely to interact again. The shorter shadow of the future in seminomadic communities may diminish the domains of attraction for stable reciprocity among nonkin (McElreath in press).

I used a narrative question (Table 2) with open discussion afterward to estimate Sangu attitudes in this domain, as well as participant observations in each community to validate the interview results. I recorded informant responses and later coded them into the categories "friend," indicating a preference for friends, and "kin," indicating a preference for kin.

The costs and accuracy of learning are likely low and high, respectively, in this domain. In the daily lives of individuals, there are many opportunities for interactions with friends and kin, and the consequences of relying on friends or kin are often easy to appreciate. Indeed, some (Cosmides 1989) argue that people are prepared by natural selection to learn when reciprocity pays and when it does not. Prepared learning of this kind would direct an individual's attention to relevant cues and reduce the search costs and increase the accuracy of individual learning. For these reasons, I code this domain into cell 1 of Table 1 (low costs, high accuracy). In some cases, costs of mistakes in this domain may be substantial. For example, relying on fair-weather friends for mutual defense could prove deadly. In terms of the model, this means b is large relative to behaving nonoptimally. As c is scaled relative to b, this increases the range of conditions that favor individual learning. The predictions derived before thus suggest that migrants between regions should change their attitudes when they change their household economies, since the reinforcements and cues relevant in this domain are tied to different modes of subsistence.

Domain 2: Respect for Elders

Pastoralist societies were found to emphasize more command and control by elders, as well as respect for them, than farmer societies (Edgerton 1971). One hypothesis that might explain this observation is that wealth is inherited more in traditional East African pastoralist settings than among farmers. In most ethnographically known pastoralist societies, sons inherit part of their father's herds, either directly or through their mother's portion. Sons also rely on parents' herds as sources of substantial bridewealth, which is paid to a potential bride's father in nearly all East African pastoralist groups in exchange for the right to marry her. Under this somewhat cynical hypothesis, individuals in pastoralists groups simply cannot afford to offend, alienate, or fall into the disfavor of elders, who control their future livelihoods (Edgerton 1971; Goldschmidt 1976).

In contrast, many agricultural economies in East Africa have only recently experienced the crowding necessary to drive land inheritance. For example, among the Nyakyusa of Southwest Tanzania (Wilson 1979), young men actually founded their own villages, inheriting little if any land and livelihood from their parents. And while bridewealth existed and exists in agricultural communities as well, it is usually much lower and frequently subject to institutions like credit. For these reasons, there is less reason in farming communities to obey, please, or fear elders, for they control less of one's future livelihood. The Nyakyusa are an usually extreme case, but the fact that especially valuable land, like that which existed in old craters, was the only land passed down to children (Wilson 1979) suggests a general importance of land scarcity and value in determining inheritance practices—especially since the Nyakyusa themselves emphasized eminently moral reasons for the lack of it.

I used a narrative (Table 2) with following discussion to estimate attitudes in this domain, as well as observation of interactions in each community. I coded responses into the categories "whenever," indicating that elders should not be feared or especially respected, and "never," indicating that elders should often or always be especially respected.

Like the previous domain, daily life provides opportunities for individuals to experience the consequences of behaving in different ways towards elders, and so individual learning can effectively adjust attitudes. There are two ways to regard costs and accuracy in this domain. First, since inheritance of wealth is a much rarer life-history event than reciprocity with friends or kin, even if individuals attend to events in the lives of their friends and family, they will have much less information in this domain than in the previous. Thus while the costs of collecting relevant information and evaluating it are likely rather low, the relative lack of information implies that learning in this domain may be much less accurate than in the previous. Second, it is probable that individuals can quickly and intuitively appreciate the power elders hold over them without much experience. If few factors other than an individual's behavior determine elders' treatment of them, then this may be the case. It is difficult to decide based on first principles which argument is more important. In the first case, I place this domain in cell 3 (low costs, low accuracy). Somehow, predicting how elders will behave seems to remind one of predicting the weather. Individual learning may be influential here but less so than in cell 1. In the second case, however, I place this domain in cell 1 (low cost, high accuracy) with the previous domain, where individual learning will strongly influence variation, and so individual economy, rather than community, will be most influential.

Domain 3: Prevalence of Witchcraft

A number of ethnographers have commented on the high prevalence of witchcraft beliefs in farmer communities, relative to pastoralist communities (see Baxter 1972). The source of this difference is unclear and there is certainly no favored theory. However, a classic and influential theory of magic (Gmelch and Felson 1980; Gmlech 1999; Malinowski 1922) suggests that people resort to magical beliefs in cases in which they have less control over outcomes. Because it is plausible that farmers experience or perceive less control over their crops than pastoralists do over their herds, Malinowski's conjecture may apply in this case as well.

Another idea, suggested by Edgerton (1971), is that witchcraft is more common among farmers as a result of repressed conflict, because it is easy for a farmer to recognize the importance of not fighting with a permanent neighbor. This repressed conflict, often aired openly among mobile pastoralists, then leads to lingering suspicion, because disagreements are rarely fully resolved. When something unfortunate happens, the still-angry neighbor becomes the prime suspect. This hypothesis is not mutually exclusive with Malinowski's, of course, and both may explain some fraction of the observed difference between farmers and pastoralists.

I used two narratives, in a between-subjects design, to get at both the overall belief in witchcraft and the tendency to nominate witchcraft as a cause of misfortune. I coded responses into two categories: (1) "other misfortune," indicating that the informant offered an explanation other than witchcraft, and (2) "witchcraft," indicating that the informant nominated magical foul play or the use of traditional medicines as the cause of the misfortune. Here I report only the first narrative (as it appears in Table 2), as the second addressed more the precise nature witchcraft beliefs than how differences are maintained.

This domain is quite different from the previous two in that costs of learning are possibly prohibitively high and accuracy is certainly low. Suppose, for example, trying to discover why a particular herd animal died. One could nominate many hypotheses: disease, angry spirits, poisonous snake bite, poisoned feed, will of God, or many others. Deciding among them may be very difficult, even for a team of veterinarians, partly because many variables pertain in such cases and simply controlling for noise in measurement is a giant epidemiological challenge. Consider how difficult it has proven to be to verify that many commonly used medicines like Vitamin C and Zinc actually do anything to prevent or treat colds. And as E. E. Evans-Pritchard (1937) noted, the existence of a nonmagical cause is not a complete explanation. If it is snakebite, why was the snake there this time? The cow had been at the water hole many times before but was never bitten. Answers to successive questions like these are much harder to come by in this domain than in either the valuation of friends or respect for elders.

This is perhaps the principle behind the cross-cultural regularity that spirits can become invisible at will; stories about visible magical beings are quickly disproved by individual experience. Because individual learning is highly ineffective for evaluating claims about supernatural phenomena, social learning can be a very powerful force in such domains. This is particularly true because humans may be designed to respond to difficult empirical circumstances with imitation, as I argued in presenting the model. For these reasons, I place this third domain in cell 4 (Table 1: high cost, low accuracy). Collecting the information necessary to evaluate competing hypotheses is expensive and time consuming and usually beyond the capabilities of an individual. Likewise, even with lots of information, some decisions are simply very difficult to make on the basis of empirical information, because so many factors influence the outcome. In this domain, we should expect that individual reinforcements, arising from personal economic arrangements, will be less influential than communities, as social learning might lead migrants to adopt the beliefs and behaviors of a majority of pastoralists, even if they are not themselves pastoralists, and vice versa.

Results

I analyzed the coded responses against individual economic variables and residence in two ways. First, I present overall differences in the frequency of responses, in order to characterize the variation between the farming and pastoralist regions. Together with these overall interview differences, I summarize informal observations gathered while living in the communities, in order to provide an understanding of how the interview responses may be reflected in other behavior. Then I present parameter estimates for logistic models predicting individual responses. These estimates indicate the influences of community residence, household economy, and previous household economy on individual attitudes, controlling for individual characteristics like gender and age.

Overall Regional Differences

General trends agree with previous findings, once household production is taken into account: 50 percent (22/44) of individuals living in the farming region said they valued friends more than kin. However, 56 percent (28/50) of individuals living in the pastoralist region said the same. This seems, if anything, the opposite of previous findings. However, classifying individuals by presence or absence of more than five cattle shows a clear difference. Sixty-four percent (35/55) of individuals in households without more than five cattle responded "friends" (see Table 2), while 38 percent (15/39) of individuals in pastoralist households responded the same (Fisher's exact one-tailed, $p = 0.0137$).

The differences in the remaining two domains are much clearer. Thirty-one percent (15/48) of farming region residents claimed that a man should never tell his father he (the father) is wrong, while 50 percent (25/50) of pastoralists said the same (Fisher's exact one-tailed, $p = 0.046$). Seventy-five percent (15/20) of individuals in the farming region responded "witchcraft" (see Table 2) when asked about ambiguous misfortune, while only 27 percent (8/30) of pastoralist region individuals mentioned witchcraft at all (Fisher's exact one-tailed, $p = 0.0009$).

Observational Evidence

Pastoralist individuals with whom I spent significant amounts of time had few friends who were not close kin. One went as far as to tell me that he would rather have "one brother than ten friends" (conversation with author, July 1997). Social life in the pastoralist regions revolved around kinship in a way it did not in the farming areas. This was not a result of the absence of non-kin to befriend: While kin tend to live close to one another, individuals also meet non-kin on most days. Also, the fact that kin live nearby one another is probably an endogenous outcome of preferences, not a cause. If pastoralists preferred friends to kin, they would presumably move households nearer to their friends. Informants in the farming region instead spent most of their social time with unrelated individuals. Several individuals, when asked about the roles of friends and kin in their lives, complained immediately about the costs of having a large family, most notably that one is obligated to share with all of them, which makes one perpetually poor.

No difference was more obvious, both to the Sangu and myself, than how young people acted towards elders in the two regions. Young adults in the farming villages gave the common respectful greetings to elders, but these greetings were often abbreviated ("*shoo*" or "*shmoo*" instead of "*shikamoo*"), which was representative of a general lack of formality in their interactions. In the interviews and later conversations, some indicated that elders in fact had no power over their lives. Several suggested that a young man should notify his father of a mistake because of an egalitarian ethic: "The father would do the same to his son" (conversation with author, July 1997). Elders themselves in the farming areas held the same opinions. Young people in the pastoralist region, however, treated elders with much more respect and ritualized interaction. A young person in the herding region seemed incapable of passing an elder member of the community without stopping, bowing slightly, casting one's eyes downward in the universal human display of subordination (Henrich and Gil-White 2001), and greeting him. If several elders were present, the young person would perform this greeting for each of them individually before continuing onward.

Finally, it is much more difficult to practice participant observation of witchcraft. Secret and locally abhorred activities are naturally the last frontier of ethnographic research. Even court records, which might provide estimates of relative prevalence of witchcraft cases, were difficult to come by, as court clerks told me that they were asked not to record such cases. One way of corroborating the interview responses, however, is by asking traditional healers *(waganga)* to list the principle reasons clients approached them. While the frequency of healers is difficult to interpret because of the substantial differences in population density between the two regions, the ranking of problems that drove the business of traditional healers should not be biased in any similar way. I interviewed two healers in the pastoralist region and three in the farmer region, asking them to list freely the most common medicines they prepared and their purposes. In the top five for both pastoralists, there were two medicines that were used only to cure curses placed by witches *(wachawi)*. The others included love potions, good-luck charms, and beauty elixirs. In contrast, for the farmers, nearly all of the medicines commonly named were used to cure curses and poisons; though one farming-region healer also listed love potions.

Logistic Model Estimates

I analyzed the results further by fitting the coded responses in each domain to a logistic model with six predictors: (1) greater than five cattle in household (binary); (2) greater than five cattle in father's household (binary); (3) community (binary); (4) age (years); (5) education (years); and (6) gender (binary). Cattle holdings were made into discrete categories because the distribution is very nonnormal, having a mode at zero. I split the sample at five cattle because this is about the number after which grazing, rather than

TABLE 3. Logistic model estimates for full models of six predictors against each domain. The estimate is followed by the odds-ratio, in parentheses, in each case. The odds-ratio here is the increased odds of observing the dependent variable when the predictor goes from its minimum to its maximum value. The standard error of each estimate appears below the significance level.

	Dependent variable (See Table 2 for meanings of codes)					
	Friends/kin "friends" N = 89		Respect for elders "never" N = 93		Witchcraft "other misfortune" N = 49	
Predictor	Estimate (Odds)	Sig (SE)	Estimate (Odds)	Sig (SE)	Estimate (Odds)	Sig (SE)
>5 cattle in household	–1.2455 (0.2878)	0.0176 (0.5417)	1.3249 (3.7620)	0.0101 (0.5260)	–0.5549 (0.5741)	0.5699 (0.9991)
> 5 cattle in father's household	–0.3146 (0.7301)	0.6371 (0.6691)	1.1629 (3.1993)	0.0626 (0.6362)	0.2174 (1.2428)	0.7743 (0.7572)
community (herders = 1)	0.3178 (1.8880)	0.3803 (0.3674)	–0.1412 (0.7540)	0.6818 (0.3463)	1.6861 (29.1411)	0.0036 (0.6783)
age (19–76 years)	–0.0468 (0.0455)	0.0521 (0.0250)	–0.0079 (0.5918)	0.7500 (0.0250)	0.0762 (153.1885)	0.0274 (0.0384)
education (0–11 years)	–0.1588 (0.1744)	0.1488 (0.1128)	–0.0203 (0.7996)	0.8578 (0.1134)	0.3862 (69.9600)	0.0345 (0.1993)
gender (female = 1)	–0.2086 (0.6588)	0.4390 (0.2702)	–0.1143 (0.7957)	0.6679 (0.2665)	0.4240 (2.3348)	0.2905 (0.4164)

corralled feeding, becomes mandatory. I omit acres under cultivation from these analyses because they in fact do not differentiate farmers from pastoralists; Sangu pastoralists farm at least as much as pure farmers do but also practice transhumant grazing. Plausible interactions for these models are not significant and do not change the characteristic results. For these reasons, I do not report them. The model estimates with six predictors are presented in Table 3.

After estimating the full models described above, I found the subset of the six predictors that lead to the optimal model, from a likelihood perspective. Adding an additional parameter to a model will almost always improve its fit to data. However, the improvement in fit may simply result from adding another degree of freedom, not from capturing a genuine causal influence. For each domain, I found the smallest set of parameters for which adding any additional parameter did not result in an increase in fit beyond that expected by chance, by a standard likelihood-ratio test (see Hilborn and Mangel 1997 for an accessible introduction). These reduced models provide more accurate estimates of the key parameters, and so I present them in addition to the full models in Table 4.

The best predictors of responses in domain 1 (preference for friends) are cattle in household and years of age. Living in a household with cattle, controlling for the presence or absence of cattle in natal household, led to a 71 percent reduction in odds of responding "friends." The kind of household an individual was raised in had a much smaller effect. The best estimate of its influence is a 27 percent reduction in odds, and this average has a very high variance. The effect of age was to reduce the odds of responding "friends": Each year of age produced a 5 percent reduction on average. While this effect is weak between individuals of adjacent age groups, across the lifespan, its effect is tremendous. It is unclear whether this estimate reflects changing attitudes with age or, rather, a cohort effect in which the value of kin has changed over time. Finally, the best estimate for the effect of community residence, controlling for household economy, is a 190 percent (approximately double) increase in odds of answering "friends" when switching between farming and pastoralist regions, on average. However, this estimate is highly variable. In the reduced model, only cattle in household and cattle in father's household are retained. The direction of the effect of father's household

TABLE 4. Logistic model estimates for optimal fit models in each domain. These models were found by likelihood-ratio criteria. Missing predictors did not improve the fit sufficiently beyond what is expected by chance when adding more parameters to the model.

	Dependent variable (See Table 2 for meaning of codes)					
	Friends/kin "friends" N = 90		Respect for elders "never" N = 94		Witchcraft "other misfortune" N = 50	
Predictor	Estimate (Odds)	Sig (SE)	Estimate (Odds)	Sig (SE)	Estimate (Odds)	Sig (SE)
>5 cattle in household	–1.2813 (0.2777)	0.0049 (0.4692)	1.1342 (3.1086)	0.0137 (0.4641)	–	–
>5 cattle in father's household	0.3822 (1.4655)	0.4112 (0.4690)	1.1579 (3.1833)	0.0113 (0.4618)	–	–
community (herders = 1)	–	–	–	–	1.0551 (8.2500)	0.0006 (0.3306)

economy switches, and it is still highly variable, suggesting its overall effect is small or inconsistent, even though it is retained in the optimal model. No plausible interaction, for example (community) × (father's cattle), appears to reduce the variance in this estimate.

Domain 2 (respect for elders) responses are best predicted by household economy and father's household economy. All other estimates are negligible. Living in a pastoralist household increases the odds of responding "never" (see Table 2) by 376 percent on average. Having grown up in a pastoralist household similar increases the odds of answering "never" by 320 percent, controlling for current household type. Community residence, controlling for household economy itself, has little effect on responses. The optimal model drops all the other predictors and produces very similar estimates for the effects of both current and natal household economy, 311 percent and 318 percent increases, respectively.

The third domain, witchcraft, shows strong effects of community residence, education and age, but smaller effects of household type. Controlling for household economy, switching to the pastoralist region resulted in a 29 fold increase in the odds of answering "other misfortune" to the witchcraft narrative in Table 2. Cattle in own household and in father's household had much smaller and variable effects. Education had the easily understood effect of decreasing likelihood of mentioning witchcraft; each year of education, on average, produced a 147 percent increase in the odds of not mentioning witchcraft. Each year of age reduced the odds of mentioning witchcraft by eight percent on average. The optimal model, however, drops the education and age effects, retaining only community of residence. The loss of education and age likely results from the smaller sample size in this domain (the sample was split among the two different witchcraft narratives). The estimate of the community effect in the reduced model is an eightfold (800 percent) increase in the odds of not mentioning witchcraft, when switching to the pastoralist community. Adding education and age back in again increases the magnitude of this estimate, but education and age themselves are not retained in the optimal model.

DISCUSSION

The results of the model estimations suggest that individual economic exposure, either own household type or the type of household one grew up in, most strongly influence the dispositions of Sangu individuals in the first two domains, preference for friends and respect for elders. In the third domain, witchcraft, individual economic exposure has little effect. Instead, community residence is the best predictor of attitudes, regardless of household type or the type of economy an individual experienced previously.

Domain 1: Preferences for Friends and Kin

I argued earlier that, in the first domain (preference for friends, Table 2), we should expect individual learning to have a strong influence, because costs are likely quite low and accuracy high (cell 1, Table 1). This leads to the prediction that individual experience should be a powerful determinant of attitudes in this domain.

The analyses above support this prediction by showing that current household economy is the only consistently strong predictor of the overall pattern of regional differences. Previous household economy has an effect, but it is inconsistent in direction in the full and optimal models. This suggests that recent and immediate circumstances most strongly influence attitudes, again consistent with the prediction that individual learning should be common and important in this domain.

Domain 2: Respect for Elders

For the second domain (respect for elders, Table 2), it is harder to produce a clear prediction. It is plausible that individuals can easily appreciate the negative consequences of disrespecting elders in pastoralist contexts, and so individual learning might be a powerful determinant of adult attitudes (cell 1, Table 1), provided that not too many factors other than behavior towards elders influence their treatment of younger individuals. However, the events in which the consequences might manifest are probably rarer than those in the first domain (preference for friends), and consequently social learning may have a stronger influence here, because of the inability for individuals to acquire enough information to make an accurate judgment (cell 3, Table 1).

The results of the analyses above are consistent with a strong influence of individual learning but suggest either a complementary influence of social learning or a developmental period during which individual experience is more determining of adult attitudes. The strong and consistent effect of current household economy is consistent with individual learning, as in the previous domain. However, the equally strong and consistent effect of previous household economy, regardless of current circumstances or community of residence, implies that recent experience is not as important as in domain 1.

Domain 3: Witchcraft

Domain 3 (witchcraft, Table 2) is likely to show the weakest influence of individual learning: Costs of acquiring effective information to make an accurate judgment are prohibitively high, and accuracy is quite low even when considerable effort is put forward to acquire relevant information (cell 4, Table 1). The influence of social learning should be very strong in this domain, reducing the influence of individual economic circumstances and increasing the influence of cultural surroundings.

The parameter estimates above are consistent with this view of the domain. Neither of the individual household economic predictors was important in either the full or reduced model. Instead, community residence was the best predictor of responses, regardless of current or previous circumstances, and has a very large effect, even in the optimal model, in which the plausibly real effects of education and age are removed and diminish the estimate greatly.

Farmers living in pastoralist communities experience very different reinforcements than pastoralists in these regions but still regard witchcraft to be less prevalent. This may be because they do not trust their own judgments about the causes of misfortune and instead are readily influenced by the ambient judgments of the majority around them. An opposite argument might hold for pastoralists living in the farming communities: While they experience very different incentives and reinforcements than their neighbors, they have good reason not to trust their own judgments about the causes of misfortune and so are easily swayed by local beliefs.

GENERAL DISCUSSION

Observational studies are plagued by a general lack of control. Even in natural experiments of the kind I have exploited here, in which similar individuals are divided by an ecological gradient, it is impossible to control for or measure even a modest proportion of the possible confounding factors. For definitive tests of hypotheses, highly controlled experiments are usually to be preferred. Nevertheless, observational studies are essential because the social sciences are presumably interested in explaining behavior in the "wild," not in the laboratory. The true test of any theory is its ability to make useful predictions about naturally occurring phenomena. Even so, it is certainly useful to mention and attempt to address a few alternative explanations and potentially confounding variables in this study. In some cases, I do not believe the current data are adequate to address the alternatives, yet it is still important to spell out what form these alternatives might take.

Is the Witchcraft Result a Product of Differential Expertise?

The result that witchcraft was mentioned less often among pastoralists could result from the fact that I asked about the death of cattle, and members of the pastoralist community simply know more about cattle and their possible causes of death. To deal with this possibility, I asked non-pastoralist members of the farming community to name potential causes of the deaths, other than witchcraft. Farmers had no trouble naming several of the most commonly named causes among pastoralists. The farmers nevertheless insisted that I had asked what someone would first think caused the deaths, and this would be witchcraft *(uchawi)*. Because farmers seem to know many of the causes nominated by pastoralists, it seems unlikely that differential experience with cattle or exposure to discussion of the causes of livestock mortality can explain the variation between the communities.

Pastoralists in the farming region certainly know additional causes but, nevertheless, were still more likely than their peers in the pastoralist region to nominate witchcraft as the cause of death, even if this is a result of exposure to ambient discussion by misinformed farmers who know little of cow biology. This, however, is not inconsistent with the interpretation I have made of the results: Even experienced pastoralists may be swayed by the poorly informed chatter of many neighboring farmers, because diagnosing the cause of an animal death is difficult and so social learning seems to them a better strategy. As I argued in presenting the model, on average across domains in which costs of learning are high and accuracy is low, individuals who employ more pure social learning will be more successful in terms of survival and reproduction. Thus the strategy in this case may be an adaptation, even if we are prepared to conclude that it leads to sometimes maladaptive consequences. As countless biologists have pointed out, success is not necessarily nor even usually the right criterion for judging an adaptive strategy.

Is Frequency Dependence Important in These Domains?

Frequency dependence arises when the optimal behavior is partly determined by what other individuals are doing. In the model I proposed, I assumed payoffs were not frequency dependent, and in the analyses, I assumed the same. If however the payoffs in domain 3 are positively frequency dependent, such that individuals are better off believing and behaving as the majority, and those in domains 1 and 2 are not, then this might explain why I found that farmers in the majority pastoralist region thought as pastoralists and pastoralists in the majority farmer region thought like farmers. This explanation does not require any social learning at all, merely the individual recognition that payoffs are frequency dependent.

In many domains, frequency dependence is very important (Gil-White 2001; McElreath et al. 2003). When many solutions exist to problems in communication, organization, and coordination, behaving as others do is likely to be a good strategy. However, it is difficult to imagine how payoffs in supernatural explanations of this kind can be frequency dependent, in this sense. However, some anthropologists (Fessler and Haley 2003) think that perhaps most domains of culture are moralized, and that individuals with minority beliefs are punished. If this is the case in this domain, then individuals may adopt majority beliefs because doing so helps them avoid punishment, as punishment creates a kind of frequency dependence. This kind of explanation does not say that individuals necessarily only hold these beliefs publicly. If punishment of minority individuals is common enough, then people may have a general tendency to believe as the majority, without even consciously considering the possibility of punishment. It is difficult to rule out a hypotheses such as this without additional data.

Are Individuals Imitating Individuals like Themselves?

The results in the first two domains raise the question of whether social learning is in fact a strong force explaining the maintenance of cultural variation here as well. If individuals who become farmers then preferentially imitate farmers, they can come to resemble other farmers even if they make no evaluations of the costs and benefits of behavior. Albert Bandura (1977) and other social psychologists have noted that individuals tend to imitate models similar to themselves, providing the psychological reality needed for this explanation. It is difficult to exclude this idea with the available data. It raises the additional question, however, of why this economy-oriented imitation is important in the first two domains but not the third. Perhaps in the first two domains it is easy enough to understand that others with the same economic system experience the same problems, yet it is still difficult to individually acquire the best behavior.

CONCLUSION

This article illustrates the possibility of investigating the maintenance of behavioral variation in a natural population. I used a model of the evolution of social learning to explore the design properties of human social learning, generating predictions about the kinds of problems for which social learning is more likely to be used and subsequently affect patterns of variation as individuals and their ideas move between regions. Measuring attitudes for several domains that vary with ecology and economy, I have tried to estimate the influence of specific individual circumstances, manifested in household economies, and ambient culture, manifested in community residence. I found that the results are consistent with the predictions of the model.

One way to view this study is that it gives comfort to all sides of debates about the maintenance of cultural variation. The model and my interpretation of the field data suggest that both those who feel that behavioral variation is a reflection of local adaptation via individual (including prepared) learning and those who feel that cultural dynamics themselves generate variation are right, depending upon the nature of the task or problem to be solved. It is important to note however that the distinction is not between adaptive and nonadaptive features of behavioral variation. Domains in which social learning is dominant can still produce adaptive variation, as the model illustrates. Yet the development (microevolution) of cultural differences and the patterning of cultural variation will be different when different mechanisms are at work. Some theories like what I have presented in this article are needed to organize our understanding of when people employ different strategies for acquiring their beliefs and behaviors, and any theory of this kind will have similarly powerful implications.

These results are important for understanding how differences in human behavior form and persist. The theory provides intuitions about the roles of the costs and accuracies of learning in any given task or problem in determining which learning strategies we might expect individuals to employ. Knowledge of the strategies individuals are likely to employ then allows prediction of the more long-term patterns of cultural diversity, including which changes in the physical and social environments are likely to enhance or erode variation. In order to explain why cultures sometimes remain distinct and why they sometimes do not, we need some theory explaining how individuals acquire their behavior and beliefs. The strategy I have employed here is to build a model of how people mix individual and social learning. Both forms of learning can lead to the maintenance or erosion of group differences, depending on the context. For example, in domains in which costs are low, and, therefore, individual learning is powerful, the model predicts that groups living in different ecologies might remain distinct regardless of the flow of persons between the groups. In a uniform ecology, however, a domain in which costs are high and social learning is important might maintain group differences despite the absence of any meaningful ecological variation. Also, as societies change, the qualities of information available sometimes change, and this alters the accuracy of individual learning. Thus as patterns of social interaction and technology change, learning strategies may change as well, leading to different patterns of behavioral variation and rates of behavior evolution.

These kinds of predictions are a beginning, but obviously many elaborations are needed to deal with additional complexities of human behavioral variation. How important is imitation of parents, peers, or other individuals? Some think that vertical transmission is especially important to human social transmission (Hewlett and Cavalli-Sforza 1986), but the evidence is weak, and others even argue that parents are largely unimportant as models (Harris 1998). Conformity (Boyd and Richerson 1985; Henrich and Boyd 1998) provides another mechanism for maintaining group variation. But there is little good evidence about how conformist transmission functions. How strong is conformity, and how much mixing is needed to overcome its influence? It is probably also true that people employ different learning strategies at different times in their lives (Henrich and Gil-White 2001), and this chronological variation in strategy will also have consequences for understanding the formation and maintenance of behavior differences. At this point, there is far more good theory than good data, but if anthropology addresses itself to this body of theory, a great deal of progress will be made.

Richard McElreath Department of Anthropology, and Graduate Groups in Ecology and Animal Behavior, University of California, Davis, 1 Shields Avenue, Davis, CA, 95616

NOTES

Acknowledgments. I thank Rob Boyd, Nick Blurton Jones, Sam Bowles, Walter Goldschmidt, Joseph Henrich, Susan Lees, Pete Richerson, Anna Simons, Annika Wallin and Tom Weisner for comments on this study and previous versions of this article. Susan Charnley helped early on by sharing Usangu contacts and information. Marehemu Ally Salum Mwalaji assisted tremendously in data collection and in my education in all things *kisangu*, and I owe him a great debt. I gratefully acknowledge funding from the

National Science Foundation and the Ford Foundation, and a generous grant from the McArthur Foundation.

REFERENCES CITED

Axelrod, Robert M., and William D. Hamilton
1981 The Evolution of Cooperation in Biological Systems. Science 211:1390–1396.
Bandura, Albert
1977 Social Learning Theory. Englewood Cliffs, NJ: Prentice-Hall.
Baxter, P. T. W.
1972 Absence Makes the Heart Grow Fonder: Some Suggestions Why Witchcraft Accusations Are Rare among East African Pastoralists. *In* Allocation of Responsibility. Pp. 163–191. Manchester: Manchester University Press.
Becker, Gary S.
1976 The Economic Approach to Human Behavior. Chicago: University of Chicago Press.
Boyd, Robert, and Peter J. Richerson
1985 Culture and the Evolutionary Process. Chicago: University of Chicago Press.
1992 Punishment Allows the Evolution of Cooperation (or Anything Else) in Sizable Groups. Ethology and Sociobiology 13(3):171–195.
Cavalli-Sforza, Luigi Luca, and Marcus W. Feldman
1981 Cultural Transmission and Evolution: A Quantitative Approach. Princeton: Princeton University Press.
Cosmides, Leda
1989 The Logic of Social Exchange: Has Natural Selection Shaped How Humans Reason? Studies with the Wason Selection Task. Cognition 31:187–276.
Durham, William H.
1991 Coevolution: Genes, Culture, and Human Diversity. Stanford: Stanford University Press.
Edgerton, Robert B.
1971 The Individual in Cultural Adaptation: A Study of Four East African Peoples. Berkeley: University of California Press.
Evans-Pritchard, Edward E.
1937 Witchcraft, Oracles and Magic among the Azande. Oxford: Clarendon Press.
Fessler, Daniel M. T., and Kevin J. Haley
2003 The Strategy of Affect: Emotions in Human Cooperation. *In* Genetic and Cultural Evolution of Cooperation. P. Hammerstein, ed. Pp. 7–36. Cambridge, MA: MIT Press.
Gil-White, Francisco J.
2001 Are Ethnic Groups Biological “Species” to the Human Brain? Essentialism in Our Cognition of Some Social Categories. Current Anthropology 42(4):515–554.
Gmlech, George
1999 Baseball Magic. *In* Conformity and Conflict. 10th edition. J. Spradley and D. W. McCurdy, eds. Allyn and Bacon.
Gmelch, George, and R. Felson
1980 Magic and Uncertainty in Modern Life. Current Anthropology 20(3):587–589.
Goldschmidt, Walter R.
1976 Culture and Behavior of the Sebei: A Study in Continuity and Adaptation. Berkeley: University of California Press.
Harris, Judith Rich
1998 The Nurture Assumption: Free Press.
Hazlewood, A., and I. Livingstone
1978 The Development Potential of the Usangu Plains of Tanzania: A Study. London: Commonwealth Secretariat.
Henrich, Joseph, and Robert Boyd
1998 The Evolution of Conformist Transmission and the Emergence of Between-Group Differences. Evolution and Human Behavior 19:215–242.
Henrich, Joseph, and Francisco J. Gil-White
2001 The Evolution of Prestige: Freely Conferred Deference as a Mechanism for Enhancing the Benefits of Cultural Transmission. Evolution and Human Behavior 22(3):165–196.
Henrich, Joseph, and Richard McElreath
2003 The Evolution of Cultural Evolution. Evolutionary Anthropology 12:123–135.
Hewlett, Barry S., and L. L. Cavalli-Sforza
1986 Cultural Transmission among Aka Pygmies. American Anthropologist 88(4):922–934.
Heyes, Cecilia, and Bennet G. Galef Jr., eds.
1996 Social Learning in Animals: The Roots of Culture. London: Academic Press.
Hilborn, Ray, and Marc Mangel
1997 The Ecological Detective: Confronting Models with Data. Princeton: Princeton University Press.
Johnson, Allen
1991 Regional Comparative Field Research. Behavior Science Research 25(1–4):3–22.
Kroeber, Alfred L., and Clyde Kluckhohn
1952 Culture: A Critical Review of the Concepts and Definitions. Vol. 42. Cambridge, MA: Museum Press.
Laland, Kevin N., and Gillian R. Brown
2002 Sense and Nonsense: Evolutionary Perspectives on Human Behaviour. Oxford: Oxford University Press.
Malinowski, Bronislaw
1922 Argonauts of the Western Pacific: An Account of Native Enterprise and Adventure in the Archipelagoes of Melanesian New Guinea. London: Routledge and KeganPaul.
McElreath, Richard
In press Social Mobility and the Strength of Norms in an African Society, the Sangu of Tanzania. *In* Cooperation, Reciprocity and Punishment: Experiments from 15 Small-Scale Societies. J. Henrich, R. Boyd, S. Bowles, C. Camerer, E. Fehr, and H. Gintis, eds. Chapter 11. Oxford: Oxford University Press.
McElreath, Richard, Robert Boyd, and Peter J. Richerson
2003 Shared Norms and the Evolution of Ethnic Markers. Current Anthropology 44(1):122–129.
Odling-Smee, F. John, Kevin N. Laland, and Marcus W. Feldman
2003 Niche Construction: The Neglected Process in Evolution. Princeton: Princeton University Press.
Pipping, Knut
1976 Land Holding in the Usangu Plain. Uppsala: Scandanavian Institute of African Studies.
Rogers, Alan R.
1988 How Much Does Biology Constrain Culture? American Anthropologist 90(4):819–831.
Shorter, Aylward
1972 Chiefship in Western Tanzania. Oxford: Clarendon Press.
Tooby, John, and Leda Cosmides
1992 The Psychological Foundations of Culture. *In* The Adapted Mind. J. H. Barkow, L. Cosmides, and J. Tooby, eds. Pp. 119–136. New York: Oxford University Press.
Trivers, Robert
1971 The Evolution of Reciprocal Altruism. Quarterly Journal of Biology 46:35–57.
Wilson, Monica Hunter
1979 For Men and Elders: Change in the Relations of Generations and of Men and Women among the Nyakyusa-Ngonde People, 1875–1971. New York: Holmes and Meier Publishers.
Wright, Marcia
1971 German Missions in Tanganyika. Oxford: Clarendon Press.

APPENDIX

Deriving the optimal mix of individual and social learning in a spatially and temporally varying environment

Assume that individual learners begin by imitating a random member of the previous generation (or a parent—since recombination between genes and culture can be ignored, it makes no difference: Boyd and Richerson 1985). They then attempt to improve on previous behavior with individual learning. This leads to a payoff to an individual learner in generation n and environment i:

$$V_{i,n}(I) = bq_{i,n-1} + b(1 - q_{i,n-1})(1 - e) - c$$
$$= b(1 - e(1 - q_{i,n-1})) - c,$$

where $q_{i,\mathrm{n}-1}$ is the frequency of behavior practiced in generation n—1 that is currently optimal. Because an individual learner can never do worse by beginning with previous behavior, assuming individual learning begins with imitation never hurts individual learning as a strategy, even when none of the previously practiced behavior is currently optimal. When $q_{i,\mathrm{n}-1} = 0$ above, the payoff is just as it would be if individual learners never imitated. Thus this model is very kind to individual learning.

Individuals may also imitate another member of the population, without attempting any individual learning, and thereby avoid the costs of trial-and-error. The expected payoff to an individual who learns purely socially in generation n and environment i is:

$$V_{i,n}(S) = bq_{i,n-1}.$$

While there is no explicit cost to social learning above, c is easily interpreted as the extra cost of individual learning.

Here we will add the additional assumption that the kind of behavior viewed as "optimal"—in terms of individual survival and reproduction—varies from time to time because of fluctuations within each environment, changing each generation with some probability v. When the environment changes, the previous state does not predict the new state, and since there are a very large number of potential behaviors, we will assume that nearly all previous behavior is rendered nonoptimal. The frequency of behavior that was practiced in generation n that is currently optimal is defined by the recursion:

$$q_{i,n} = (1 - v)(1 - m)(1 - \alpha)(1 - e(1 - q_{i,n-1})) + \alpha q_{i,n-1}$$

where α is the proportion of individuals in the population who learned socially in the previous generation. If the environment has not changed and rendered a new behavior optimal, social learners acquire optimal behavior proportional to its frequency in the previous generation. When the environment changes, however, almost none of the behaviors practiced in the previous time period are adaptive, so the expected frequency of adaptive behavior available to imitate is approximately zero. Likewise, a proportion m of the population after migration is from other environments and are unlikely to carry any optimal behavior.

The goal is to know what mix of social and purely individual learning is favored by natural selection. Obviously individuals are not only social learners or individual learners, but instead sometimes purely imitate and sometimes also learn on their own. I formalize this by letting individuals practice any mix of individual and social learning, such that they learn individually 1–α of the time and socially α of the time. It is reasonable to assume that the dynamics of learning take place on much shorter time scales than those of natural selection on genes (Boyd and Richerson 1996; Rogers 1988). In addition, it is easy to show numerically that q_i quickly reaches an equilibrium for any given mix of individual and social learning in the population. The equilibrium value of $q_i, \hat{q}_i$, occurs where $q_{i,n} = q_{i,n-1}$ and is:

$$\hat{q}_i = \hat{q} = \frac{(1 - \alpha)(1 - v)(1 - m)(1 - e)}{1 - (1 - v)(1 - m)(\alpha + (1 - \alpha)e)}.$$

As all environments i are alike in these terms, this equilibrium is the same in all environments.

The expected payoff to a mutant individual who learns socially $\alpha' = \alpha + \delta$ of the time, once the frequency of optimal behavior goes to equilibrium for a population of individuals with proportion of social learning α, is:

$$\hat{V}(\alpha') = (1 - \alpha')\{(1 - e(1 - \hat{q}))b - c\} + \alpha' b\hat{q}.$$

The evolutionarily stable (CSS) value of α, the only proportion that cannot be invaded by individuals employing any other proportion α', is found where

$$\left.\frac{\partial \hat{V}(\alpha')}{\partial \delta}\right|_{\delta=0} = 0.$$

Solving the above for α yields the evolutionary stable proportion of individual learning, $\alpha^{\hat{}*}$:

$$\alpha^* = \frac{1 - e(1 - v)(1 - m) - (1 - e)(v + (1 - v)m)b/c}{(1 - e)(1 - m)(1 - v)}$$

When $\alpha^* \geq 1$, natural selection favors only individual learning. When $\alpha^* \leq 0$, natural selection favors pure social learning, because individual learning is so costly or inaccurate that one would be better off guessing. When there are no individual learners, cultural evolution cannot track changes in the environment, and the population in fact ends up guessing.

Note that, in this model, temporal variation v is completely analogous to spatial variation. For this reason, I ignore the effects of v in the main text. There are interesting differences between the two types of variation, however (Henrich and Boyd 1998).