Chapter 3: Culture Evolves

"When a dog bites a man, that is not news," goes the journalistic aphorism, "but when a man bites a dog, that is news."[1] To many anthropologists, the claim that culture evolves will seem more like "Dog bites man" than "Man bites dog"—it may or may not be true, but it certainly is not news. In fact, the idea that culture evolves is as old as the discipline of anthropology itself. The nineteenth-century founders of anthropology, Lewis Henry Morgan and Edward Tylor,[2] thought that all societies evolved from less complex to more complex through the (in)famous stages of savagery, barbarism, and civilization. Such progressive evolutionary theories continued to be important throughout most of the twentieth century in the work of noted anthropologists like Leslie White, Marshall Sahlins, Julian Steward, and Marvin Harris. During this period, evolutionary theories became less ethnocentric and more realistic. Evolutionary stages were given less-loaded terms such as *bands, tribes, chiefdoms,* and *states,*[3] and models were developed that allowed for the effect of local ecology on the trajectories of cultural evolution.[4] Though evolutionary theories no longer dominate contemporary anthropology, though they continue to have important defenders like Robert Carneiro, Allen Johnson, and Timothy Earle.[5] The attraction of such progressive evolutionary theories is plain to see. The archaeological and historical records leave no doubt that the average human society has become larger, more productive, and more complex over the last ten thousand years. Although unilineal theories of human progress have fallen out of favor, the general trend toward greater complexity is not in doubt.[6]

However, we mean something quite different when we say culture evolves. Remember that the essential feature of Darwin's theory of evolution is population thinking. Species are populations of individuals who carry a pool of genetically acquired information through time. All of the large-scale features of life---its beautiful adaptations and its intricate historical patterns---can be explained by the events in individual lives that cause some genetic variants to spread and others to diminish. The progressive evolutionary theories debated by generations of anthropologists have almost nothing in common with this Darwinian notion of evolution. Very little of their work focuses on the processes that shape cultural variation; it is mainly descriptive. Those accounts of cultural evolution that do provide mechanisms typically focus on external causes of change. People's choices change their environment, and these changes lead to different choices. For example, a common argument is that the evolution of political and social complexity is driven by population growth—denser populations require economic intensification and facilitate political complexity, division of labor, and so on.[7] Such processes are more akin to ecological succession than evolution. In the same way that lichen colonizing a glacial moraine change the environment, making the soil suitable for grasses which in turn further change the soil, making way for shrubs, simpler societies change their environments in ways that make more complex societies necessary.

There is little doubt that such successional processes have played a role in human history. However, they are far from the whole story;[8] culture evolves. Human populations carry a pool of culturally acquired information and in order to explain why particular cultures are as they are, we need to keep track of the processes that cause some cultural variants to spread and persist while others disappear. The key is to focus on the details of individual lives. Kids imitate one another, their parents, and other adults, and both children and adults are taught by others. As children grow up they acquire cultural influences, skills, beliefs, and values which affect the way that they lead their lives, and the extent to which others imitate them in turn. Some people may marry and raise many children, while others may remain childless but achieve prestigious social positions. As these events go on year after year and generation after generation, some cultural variants thrive while others do not. Some ideas are easier to learn or remember, some values are more likely to lead to influential social roles. The Darwinian theory of cultural evolution is an account of how such processes cause populations to come to have the culture they have.

This sketch of a Darwinian theory of culture emphasizes the generic properties of different types of processes. For example, some cultural variants may be easier to learn and remember than others, and this will, all other things being equal, cause such variants to spread, a process we call biased transmission. The basic kinds of processes are the *forces* of cultural evolution, analogous to the forces of genetic evolution, selection, mutation, and drift. In any particular situation, the concrete events in the lives of real people are what really goes on.. However, by collecting similar processes together, and working out their generic properties, we build a handy conceptual tool kit that makes it easier to compare and generalize across cases. While we make no pretense that our scheme is a finished and final account, we do think that the tools in hand are useful and can be used to help understand how culture evolves.

A Darwinian account of culture does not imply that culture must be divisible into tiny, independent genelike bits that are faithfully replicated. Rather, the best evidence suggests that cultural variants are only loosely analogous to genes. Cultural transmission often does not involve high-fidelity replication; nor are cultural variants always tiny snippets of information. Nonetheless, cultural evolution is fundamentally Darwinian in its basic structure. Analogies to ordinary biological evolution are useful but only because they provide us with a handy, ready-made tool kit to use in building a theory rooted in the best social science.

Skeptics who distrust Darwinism are common, particularly in the social sciences. But, Darwinism is not inherently an individualist, adaptationist footpad sneaking into the social sciences to explain everything by genetic reductionism. Nor does it signal a return to the progressive, Eurocentric ideas of the past. A great variety of substantive theories arise when the all-important details are specified. Some models end up looking a lot like rational choice; and in others, arbitrary cultural differences can arise from the dynamics of interacting cultural elements. Some models lead to long-term directional change in which artifacts or institutions become more efficient, while others lack such trends.

Culture is (mostly) information in brains

The first step in applying population thinking to human culture is to specify the nature of the information that is being transmitted. Culture is (mostly) information stored in human brains, and gets transmitted from brain to brain by way of a variety of social learning processes.

Every human culture contains an enormous amount of information. Think about how much information must be transmitted just to maintain a spoken language. A lexicon requires something like sixty thousand associations between words and their meanings. Grammar entails a complex set of rules regulating how words are combined into sentences; and although some of these rules may arise from innate, genetically transmitted structures, clearly the rules that underlie the grammatical differences separating languages are culturally transmitted. Subsistence techniques also entail large amounts of information. For example, the !Kung San of South Africa have a very detailed knowledge of the natural history of the Kalahari Desert—so detailed in fact that the researchers who studied them were unable to judge the accuracy of much of !Kung knowledge, because it exceeded the expertise of Western biology.[9] As anyone who has ever tried to make a decent stone tool can attest, the manufacture of even the simplest implement requires lots of knowledge; more-complex technology requires even more. Imagine the instruction manual for constructing a seaworthy kayak from materials available on the north slope of Alaska. The institutions that regulate social interactions incorporate still more information. Property rights, religious custom, roles, and obligations all require a considerable amount of detailed knowledge to make them work.

The vast store of information that exists in every culture must be encoded in some material object. In societies without widespread literacy, the main objects in the environment capable of storing this information are human brains and human genes. Undoubtedly some cultural information is stored in artifacts. The designs that are used to decorate pots are stored on the pots themselves, and when young potters learn how to make pots they use old pots, not old potters, as models. In the same way, the architecture of the church may help store information about the rituals performed within. Without writing, however, artifacts can't store much information. The young potter cannot learn how to fire a pot simply by studying existing ones. Without written language, how can an artifact store the notion that Kalahari porcupines are monogamous, or the rules that govern bride-price transactions? With the advent of literacy, some important cultural information could be encoded on the pages of books.[10] Even now, however, the most important aspects of culture still tend to be those stored in our heads.

Behavior depends on skills, beliefs, values, and attitudes

Unfortunately, there is little scientific agreement about *how* information is stored by human brains. In some parts of the social sciences, especially history, people's behavior is often understood in terms of their values, desires, and beliefs. In other parts of the social sciences, the notions of values and beliefs are formalized under the "rational actor" model, in which values are represented by a "utility function," a mathematical rule that assigns a number to every state of the world that an individual might experience. Beliefs are represented as a Bayesian probability distribution that specifies the individual's subjective probability that each state of the world will occur. Individuals make choices that maximize the expected value of their utility. Many find the rational actor account of human psychology to be convincing because of its theoretical elegance; mathematicians have shown that only by maximizing expected utility can people avoid grossly irrational behavior—preferring ice cream to pickles, pickles to pizza, and pizza to ice cream, for example.

Psychologists of all stripes caution us that values and beliefs are folk psychology, culture-bound folk psychology at that,[11] and most care nothing for formal elegance and everything for empirical verisimilitude. Psychologists also believe that the brain is crucial for understanding all aspects of human behavior, from "low-level" functions such as processing visual information to "higher-level" functions such as reasoning or speech production. Since the real world of the human mind is complex and poorly understood, deep disagreements exist within psychology about how such information is stored and how it shapes behavior. Behaviorists concentrate on observable behavior and cognitive scientists speak of mental rules and representations,[12] while others deny the relevance of such entities and argue that only neurophysiological descriptions are useful.[13] It is unclear whether these pictures of the human mind can be integrated. To quote the eminent psycholinguist Ray Jackendoff,

> What is pretty much a mystery at this point is how linguistic rules and representations are neurally instantiated—that is how [the] physical structure of the brain could make

> possible the combinatorial regularities discovered by linguistic research. In fact, other than certain aspects of low-level vision, I know of no success at relating systematicities of mental representations to the details of neural architecture.[14]

A lot of progress can be made without solving these problems. However, we need some expedient agreement about what to call the information stored in people's brains. This problem is not trivial, because psychologists have deep disagreements about the nature of cognition and social learning. Adopting a terminology may mean taking sides in these controversies, something that is neither necessary nor desirable. But, we can't go on saying "information stored in people's heads"—it's just too awkward. Some authors use the term *meme* coined by the evolutionary biologist Richard Dawkins, but this connotes a discrete, faithfully transmitted genelike entity, and we have good reasons to believe that a lot of culturally transmitted information is neither discrete nor faithfully transmitted. So we will use the term *cultural variant.* We will also sometimes use the ordinary English words *idea, skill, belief, attitude,* and *value* without meaning to imply that introspection is necessarily a reliable guide to what is stored in your own brains, or that what people tell you is necessarily a reliable guide to what is stored in *their* brains. Psychologists will one day exchange the terms of folk psychology for clearly defined, scientifically reliable concepts; in the meantime we use these terms in the interests of producing readable prose.

Cultural variants are acquired by social learning

Many of the beliefs, ideas, and values that influence people's behavior are acquired from other people through social learning,[15] We will loosely say that people imitate other people, but in fact ideas get from one head to another by a variety of complex processes. Consider how you learn to tie a knot, say, a bowline. As simple as it is, almost no one invents such a clever knot; they learn

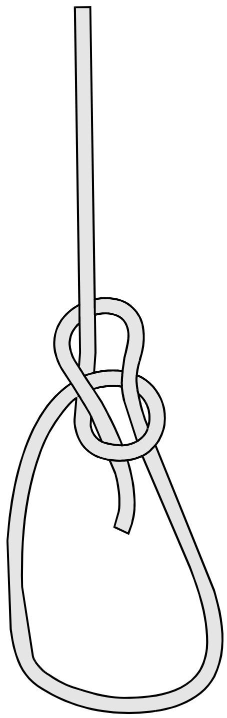

Figure 3.1. Although the bowline knot is strong and easy to untie, it can accidentally come untied.

it from others, but they do so in many different ways. Some learn by verbal instruction. Someone tells you that a bowline is a strong knot that can nonetheless be easily untied. Someone else teaches you the algorithm "The rabbit comes out the hole, up the tree, around the tree, and back down the hole." You can learn by watching somebody tying a bowline, or you might come upon an example of a bowline in a book and learn how to tie it by yourself. You can learn from us by studying the picture in figure 3.1. (Try it! It's much better than an overhand knot for many everyday tasks.) What these forms of social learning have in common is that information in one person's brain generates some behavior---some words, the act of tying a knot, or the knot itself---that gives rise to information in a second person's brain that generates a similar behavior. If we could look inside people's heads, we might find out that different individuals have different mental representations of a bowline, even when they tie it exactly the same way.

Cultural evolution is Darwinian

Now, let's see how we can use population thinking to link these facts about how culture stored and transmitted by individuals to the two central facts about cultural variation: traditions exist, and traditions change.

Consider a simple, hypothetical example inspired by Salamon's account of German and Yankee farmers. This is not a real model of cultural evolution in Illinois, rather, it is way of illustrating the logic of Darwinian methods.[16] The standard way to modularize an evolutionary problem is to think about the main events in the life cycle of an individual, divide that life cycle into stages in which only one process operates, specify the processes, develop the statistical machinery to scale up from individuals to the population, and then use this machinery to keep track of the distribution of cultural variants as the population marches through history, one generation at a time.

First, we must define the problem. What are the boundaries of the population? And, what cultural variants are present in the population? Assume that basic values about farm and family are only acquired from members of the local community, which means that we can take the community as our population. If we were interested in the evolution of some other trait, say, preferences for recorded music, the population would be different, because these preferences are strongly influenced by people outside the community. Let us also assume that there are only two variants: people have either yeoman values or entrepreneurial ones. Of course, reality is much more complicated, and we will consider how to deal with such complications later on; but for now it helps to keep things simple. We also need to decide how to represent the distribution of cultural variants in the population at any one time. Because there are only two variants, it is convenient for this purpose to keep track of the fraction of the population who hold each belief. In other situations we use other statistics to describe the distribution of beliefs.

Next, we consider what happens at each stage of the cultural "life cycle" (fig. 3.2). Here we assume that children initially acquire the beliefs of their biological parents. Children growing up in families with two parents having yeoman values acquire yeoman values; children with two entrepreneurial parents acquire entrepreneurial values; and children whose parents differ are equally likely to acquire yeoman values and entrepreneurial values. This means that transmission from parents to offspring leaves the population unchanged from one generation to the next. This

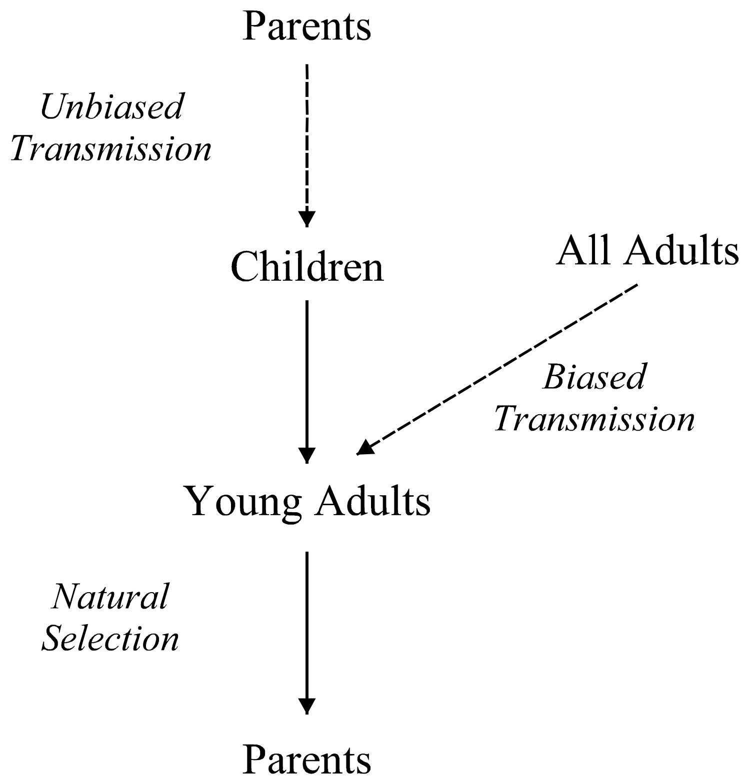

Figure 3.2. A diagram of the life cycle described in the text. Children acquire beliefs and values about farming from their parents. Then, as they grow older, their beliefs and values may also be affected by other adults. Next, as adults, they marry and choose a career. Those who abandon farming and leave the community have no further effect on the values in the community

model assumes accurate replication of cultural variants, although social learning in practice will probably introduce frequent errors.[17] The basic framework can easily be modified to allow for this possibility.

As children grow older they are exposed to people other than their parents, some of whom may cause them to modify their beliefs. Suppose that young adults get experience with other farm operations (perhaps as the result of participating in young-farmer groups like 4-H). They observe that farmers with yeoman values work longer hours and make less money, but have closer family ties than do their entrepreneurial counterparts. These observations cause some young adults to adopt new values—some switch from yeoman values to entrepreneurial ones, and some do the opposite. For most young adults a close family doesn't compensate for long days and low wages, so more of them switch from yeoman to entrepreneurial values than the reverse. This is an example of *biased* cultural transmission, which occurs when people tend to acquire some cultural variants rather than others. Biases may be innate preferences, or they may be cultural preferences acquired in an earlier episode of social learning.

Eventually, young adults grow up. Some obtain a farm and remain in the community, while others abandon farming to become mechanics, salespeople, lawyers, or academics. Salamon's data suggest that people who hold yeoman values are more likely to remain in the community. Since only adults who remain in the community influence the values of community members of the next generation, selective emigration, a type of natural selection of cultural variants, has the effect of increasing the proportion of the community holding yeoman values.

Finally, people get married and have children. According to Salamon, the descendants of German immigrants have about 3.3 children per family, while those descended from Yankees have only 2.6.[18] Suppose that this difference in family size results from the same belief system that causes differences in farm management and inheritance patterns. Since children initially acquire their values from their parents, this means that differential reproduction also leads to the spread of yeoman values in the community. This process is another form of natural selection, and rather strong selection at that.

Now let's use this model to explain the why cultural differences persist. So far we have seen how various processes lead to cultural continuity and change within a single generation. To explain long term persistence, we iterate the model from generation to generation to determine what happens over time.

The ancestors of the Yankees and Germans of Salamon's study came to Illinois with them different values which led to significant differences in behavior even though they farmed on similar soils and faced the same technical and economic constraints. In the simplified world of our model, this means that the net effect of all the social learning processes operating in each population is to leave each population more or less unchanged. If yeoman values are common in one generation, then they will be common in the next. If entrepreneurial values are common, they will remain common.

"Cultural inertia" can arise in two ways. It can arise from a tendency to conform to the beliefs of the majority. However, in the current model, the most natural explanation is a combination of unbiased sampling and faithful copying. You can think of children as being exposed to a sample of two of the cultural variants of the previous generation. Sometimes both parents hold entrepreneurial vales, sometimes both hold yeoman values, and sometimes parents differ. As long as holding yeoman values doesn't have too big an effect on family size, these samples will be representative of the population from which they were drawn, meaning that the probability that a parent holds yeoman values is approximately the same as the frequency of yeoman values in the population. Then, as long as the cultural learning process is accurate and unbiased, the probability that a child acquires yeoman values will also be approximately the same as their frequency in the population of parents—transmission from parents to offspring won't change the cultural composition of the population. The same goes for social learning by young adults. Again, they are exposed to a sample of adults from the previous generation. If the sample is representative of the population, and if young adults are not strongly predisposed to acquire entrepreneurial values, transmission will lead to little change.

We also want to explain how cultures change. In the present case, there are three possibilities. One possibility is that the effect of biased transmission is very strong—almost everyone who starts out with yeoman values switches to entrepreneurial ones, and almost everyone who starts with entrepreneurial values retains them. Then entrepreneurial values will spread in the community, because people are predisposed to choose such values. Second, biased transmission could be relatively weak—some people switch from one set of values to another, but most people retain the values that they learned from their parents. Then yeoman values will spread, because people with such values are more likely to stay in the community and because they have larger families. This is what actually seems to be happening in the communities that Salamon studied. Third, the community might settle down to a stable mix of the two types.

The forces of cultural evolution

We call the processes that cause the culture to change *forces* of cultural evolution. We divide the evolving system into two parts. One is the "inertial" part—the processes that tend to keep the population the same from one time period to the next. In this model cultural inertia comes from unbiased sampling and faithful copying of models. The other part consists of the forces—the processes that cause changes in the numbers of different types of cultural variants in the population. These processes overcome the inertia and generate evolutionary change.[19]

In our stripped-down version of the lives of German and Yankee farmers, two forces are at work. Biased transmission causes entrepreneurial values to increase, and natural selection causes yeoman values to increase. These two processes exemplify two distinct classes of forces. Transmission biases are forces that arise because people's psychology makes them more likely to adopt some beliefs rather than others. Natural selection is a force that results from what happens to people who hold different cultural variants. We focus on biased transmission and natural selection here as a device to introduce the logic underlying our models of cultural evolution, and in subsequent chapters we extend our analysis to include the other forces introduced in table 3.1.

Table 3.1 A list of cultural evolutionary forces discussed in this book

Random forces

Cultural mutation. Effects due to random individual level processes like misremembering an item of culture.

Cultural drift. Effects due to statistical anomalies in small populations. For example, in simple societies some skills, like boat-building, may be practiced by a few specialists. If all the specialists in a particular generation happen, by chance, to die young or to have personalities that discourage apprentices, boat-building will die out.

Decision-making forces

Guided variation. Non-random changes in cultural variants by individuals that are subsequently transmitted. Results from transformations during social learning, or learning, invention, or adaptive modification of cultural variants.

Biased transmission

Content-based (or direct) bias. Individual are more likely to learn or remember some cultural variants based on the their content. Content based bias can result from calculation of costs and benefits associated with alternative variants, or because the structure of cognition makes some variants easier to learn or remember.

Frequency-based bias. The use of the commonness or rarity of a cultural variant as a basis for choice. For example, the most advantageous variant is often likely to be the commonest. If so, a conformity bias is an easy way to acquire the correct variant.

Model-based bias. Choice of trait based on the observable attributes of the individuals who exhibit the trait. Plausible model-based biases include a predisposition to imitate successful or prestigious individuals, and a predisposition to imitate individuals similar to oneself.

Natural selection.

Changes in the cultural composition of a population due to the effects of holding one cultural variant rather than others. The natural selection of cultural variants can occur at individual or group levels.

Biased transmission

Biased cultural transmission occurs when people preferentially adopt some cultural variants rather than others. Think of it as comparison shopping. People are exposed to alternative ideas or values and then choose among them (although the choice may not be a conscious one).[20] The diffusion of innovations provides a fund of well-studied examples of how biased transmission works. This body of work was pioneered by a landmark study by sociologists Bryce Ryan and Neal Gross of the spread of hybrid corn (maize) in two Iowa farm communities by in the early 1940s. Following their lead, thousands of case studies of the diffusion of innovations have been published.[21] These studies indicate that in both traditional and contemporary societies, innovations often spread as the result of personal contact. People adopt an innovation like hybrid maize after observing the behavior of friends and neighbors who have already adopted the innovation. Once they have observed the innovation firsthand, their decision about whether to adopt the innovation is strongly affected by the perceived utilitarian advantage of the new crop. Is the hybrid seed more resistant to disease? Is there a ready market for the new crop? If so, people will tend to adopt the new crop and the innovation will spread.[22] The decision to adopt a new idea, crop, or any other cultural variant may also be affected by the number or prestige of the people who have already adopted it leading to varieties of biased transmission that we will consider in detail in chapter 4.

Because biased transmission results from the (not necessarily conscious) comparison of alternative variants, the resulting rate of cultural change depends on the variability in the population. Initially, innovations spread slowly because few people practice them, and so few other people are in a position to observe the innovation and compare it with their existing behavior. As the innovation becomes more common, more people are exposed to it and can

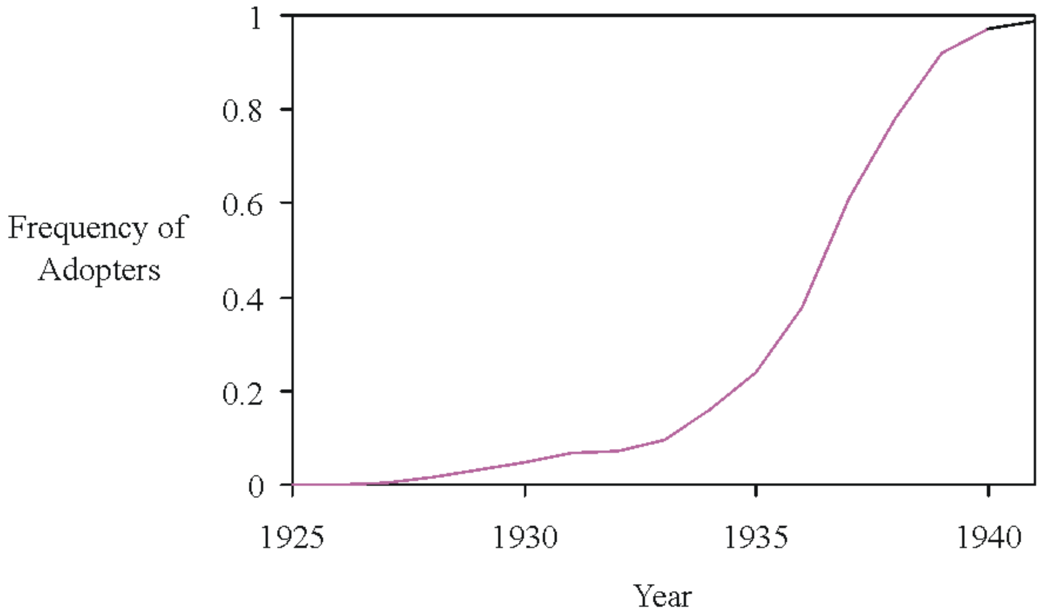

Figure 3.3 The spread of hybrid corn (maize) in the American Midwest. The S-shaped trajectory is consistent with the hypothesis that biased transmission played an important role in this diffusion. Initially, the rate of spread was low because very few farmers had adopted the hybrid corn and therefore the chance of encountering a farmer who used the hybrid seed was also low. As adopters increased in frequency, the rate accelerates. As more and more farmers adopt the hybrid corn, the rate of innovation decelerates.

compare it with other behaviors, and the rate of adoption of the innovation accelerates. As the old behavior becomes rare, there are fewer people still practicing it and fewer opportunities to make the comparison, so the rate of spread of the new behavior slows. This process, which has been documented in many different cases, generates a characteristic S-shaped trajectory. like the one shown in figure 3.3.

The rate at which a population changes by biased transmission also depends on how hard it is to evaluate alternative behaviors. If a new crop variety has substantially higher yields than existing crops, then farmers will easily detect the difference. Hybrid corn had about a 20% yield advantage over traditional varieties, so its use spread rapidly. Similarly, after sweet potatoes were introduced to coastal New Guinea from the New World sometime in the 1700s, they swifly replaced other crops in the cool highlands because they performed much better typical tropical plants. This happened even though the Europeans who brought the sweet potatoes to New Guinea went no further than the coast and didn't even know that people lived in the highlands until the 1930s.[23] However, the benefits of many other desirable traits may be much harder to detect. The practice of boiling drinking water substantially reduces infant mortality from diarrhea. Nonetheless, the practice may fail to spread, because the effects of boiling water are difficult to discern. There are other ways of getting diarrhea, and people can't see the microbes in the water. People who believe that disease is caused by magic may find it hard to believe that boiling drinking water is useful. Figuring out which variant is best is often hard even if they have very different payoffs. Traits whose beneficial effects only become apparent over time are especially difficult to evaluate.

Biased transmission doesn't always result from an attempt to evaluate alternative cultural variants according to cultural standards or rules. Biases are often caused by universal characteristics of human cognition or perception. For instance, many linguists believe that some linguistic features are "marked," meaning that they are harder to produce and perceive than alternative unmarked features. Languages that denote the subject and object of sentences with word order are less marked that languages accomplishing this function by changing the form of the noun. Such unmarked features are simpler, and accordingly appear earlier in first language acquisition. Many linguists also believe that "internal" language change (as opposed to change that results from contact between languages) typically proceeds from marked to unmarked. Such changes tend to make the language easier to produce and understand. Thus, language learners confronted with two slightly different grammatical systems will tend to adopt the less marked of the two, and in this way biased transmission can drive language change.[24] This hypothesis is somewhat controversial, but if it turns out to be true, it will provide a good example of how biases may arise from the workings of human psychology.

Biased transmission depends on learning rules

The strength and direction of biased transmission always depend on what is going on in the minds of imitators. The explanation for the increase in the frequency of entrepreneurial values in rural Illinois lies in the values of young adults. Why do they value cash and comfort over family? In some cases, values may result from universal human propensities—desires for wealth, comfort, and control over your life are likely built into human psychology. In other cases, values may stem from other cultural variants—cash and comfort might win in contemporary Illinois, but family loyalty win in rural China.

Anthropologist William Durham distinguishes between genetically acquired learning rules, which he calls "primary value selection," and culturally acquired learning rules, which he refers to as "secondary value selection."[25] The rules that underlie change in the way words are pronounced (phonology, in linguistic jargon) provide a good example of this distinction. To a first approximation, the pronunciation of vowels can be represented in a two-dimensional space which represents the vertical and horizontal position of the tongue. Ample evidence from many different languages shows that pronunciation evolves so that the distance between vowels in this space is maximized. Presumably, people subconsciously prefer widely spaced vowels because they facilitate both pronunciation and understanding.[26] Young people who are establishing their dialect listen to the pronunciation of others and tend to adopt speech variants of people whose vowels are most evenly spaced. That this process has been documented in a wide range of different languages suggests that the preference for evenly spaced vowels is what Durham would call a primary value.

Language change also provides examples of secondary values. When people speaking different languages come into contact, all kinds of linguistic variants can diffuse from one language to the other. The rate at which this occurs depends on how similar the languages are. When languages are similar, people hear a new form, find it understandable, and then can incorporate into their own language. If languages are very different, its harder to learn foreign words or grammatical forms, and borrowing is inhibited. Thus, the attractiveness of a new form depends on the language that you and your community already speak, which is an example of what Durham labels secondary values.

The relative importance of primary and secondary values selection is controversial. Some evolutionary biologists, such as Richard Alexander, Charles Lumsden, and Edward Wilson, advocate a dominant role for primary values.[27] Durham makes a case for the importance of secondary values, although his terminology implies that secondary values derive from the primary ones. Our hunch is that primary and secondary values virtually always interact. Consider the effects of contact-induced language change. The usefulness and intelligibility of new forms is governed by the similarity of the two languages in contact. But why do people want to communicate effectively? Why don't people choose the less- rather than more-intelligible forms? Sometimes they do: think of lawyers, politicians, or sometimes, alas, scientists.[28] People may prefer gratuitously complex linguistic forms to signal that they occupy of particular social role or for similar culture-specific reasons. The reason people often do prefer less marked forms must lie in the basic nature of human psychology—people (usually) want to be understood. The difficulty of convincing people to boil their drinking water illustrates the same point. The desire to avoid unnecessary work, like gathering extra fuel to boil water, and the desire for children to thrive are likely to be primary values that have deep, genetically influenced psychological roots. Belief in the germ theory of disease creates a secondary value. The decision about whether or not to boil drinking water depends on both these primary and secondary values.

How cultural variants compete

So far we have tacitly assumed that cultural variants compete with each other.[29] farmers either hold yeoman values *or* entrepreneurial values, that people use one dialect *or* another, that they either adopt innovation *or* retain their present behavior. This either/or dichotomy is appropriate for genes, but it may not be for culture. The competition between different versions of the same gene results from the machinery of genetic replication. Every gene sits at a particular site, or locus, on a particular chromosome. For example, in a population of one thousand individuals there are two thousand chromosomes that can carry any given gene. If the number of chromosomes that carries one version of a gene increases from one generation to the next, the number of chromosomes that carries alternative versions of the same gene must decrease. Cultural replication need not have the same dichotomous character. People can learn and remember more than one variant. For example, they could know how to speak two different dialects, so a new dialect can spread through a population without other dialects declining.

We think that cultural variants compete in two related ways. First, they compete for the cognitive resources of the learner, both during the process of social learning, and afterward when

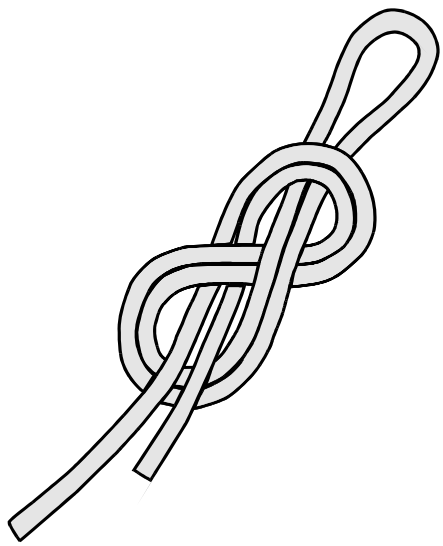

Figure 3.4. The figure-eight knot is strong, unlikely to become accidentally untied, and easy to untie if loaded, but it is a bit slow to tie.

the learner must expend some effort maintaining the variant in memory. Learning things takes time and energy that could be devoted to other valued activities, and may compete with remembering old ones. This constraint may not be very important for knowledge that is easy to acquire. For example, a bowline, a fisherman's knot, and a figure-eight knot (fig. 3.4) can all be used to tie a loop at the end of a rope, and you easily can learn them all. The amount of time that it takes you to learn a bowline doesn't prevent you from learning the others. Learning a new knot takes only a few minutes.

But, for knowledge that is more difficult to acquire, the cost of learning leads to sharp competition between variants. Mastering a new academic discipline or learning a new language requires a substantial investment of time and energy, and this may require us to choose among alternatives. Some years ago we spent a year at a German university, and we both thought it would be a good thing to learn German, but we both chose to spend the time working on this book instead. Competition between cultural variants for time and energy is diffuse compared to competition between genes at a locus. It does not necessarily lead to competition between variants that affect the same behavior; rather, it causes competition between all the variants that a person might acquire at a given time. German did not compete with French for our all-too-limited time and attention. German competed with learning historical linguistics and studying the history of technology. The diffuse competition for our time and energy seems to limit our willingness to build up big repertoires of even such simple, useful skills as tying knots.

The second and, we think, more stringent form of competition between cultural variants is for control of behavior. People learn a great deal by observing others, and if a cultural variant doesn't affect behavior, it won't be transmitted. Unlike genes, culture is a system of inheriting acquired variation. It has no analog of recessive or silent genes that do not influence phenotype—an organism's observable properties produced by the interaction of its genetic material with the environment—but are transmitted anyway. If you believe that a figure-eight ring knot is the best knot for making a loop in the end of a rope, and you always use this knot, then the people cannot learn other knots from you, even if know how to tie other knots. The competition between cultural variants will be particularly acute when they affect many aspects of a person's life. An Illinois farmer who holds yeoman values will behave differently almost every day of his life than one who holds entrepreneurial values. Protestant converts to Catholicism or Buddhism may remember all the Protestant doctrine they learned, yet they will cease being models for Protestantism.

A long-unused variant may also be forgotten. We have all experienced the distressing loss of some hard-earned skill like differentiation, playing the clarinet, or carving a parallel turn. Use it or lose it.

People also learn ideas and values through overt teaching.[30] Here the effect is more subtle. The same sorts of things that cause a cultural variant to be used will also cause it to be taught, and to be used by those who acquire it. If you believe that the figure-eight knot is the best knot because it is strong, unlikely to accidentally untie, but easily untied after being under tension, then it is likely that this will be the knot you teach to others. Even if you teach people to tie other knots, they will be more likely to use the figure-eight knot if they accept your argument about why it is best.

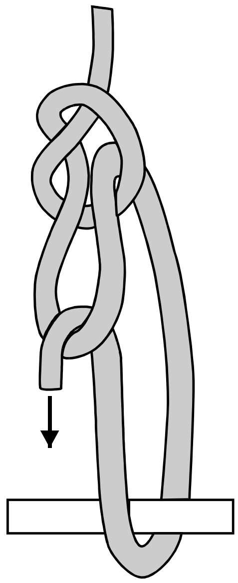

Figure 3.5. The trucker's hitch knot is useful for securing loads because of its mechanical advantage.

Competition for control of behavior is much less diffuse than competition for attention. If two variants specify different behavior in the same context, typically only one of them can control behavior. We can drive on the right or the left, but only drunks and foolish teens try both. In bilingual environments people may switch rapidly from one language to the other, even in midsentence, but word by word, or at least word fragment by word fragment, they can be speaking only one. This example also illustrates the interaction between the two forms of competition. If a trait can easily be learned, it will not matter so much that rarely affects behavior—occasional demonstrations will allow it to perist. One of us learned a rare but very useful knot, the trucker's hitch (fig. 3.5), from a single demonstration many years ago, the first and only time he has ever seen anyone else tie it. On the other hand, skills and knowledge that are only acquired over long periods of observation will be strongly affected by the amount of time that they can be observed.

Natural selection of cultural variations

The logic of natural selection applies to culturally transmitted variation every bit as much as it applies to genetic variation. For natural selection on culture to occur,

- people must vary because they have acquired different beliefs or values through social learning,
- this variation must affect people's behavior in ways that affect the probability that they

transmit their beliefs to others, and

- the total number of cultural variants that can exist in the population must be limited in some way.

Or, in other words, cultural variants must compete.

You can substitute the appropriate genetic terms in this list to recover the standard textbook account of how genes evolve by natural selection. The basic logic is identical. All other things being equal, beliefs that cause people to behave in ways that make their beliefs more likely to be transmitted will increase in frequency. If the behaviors that are shaped by the beliefs acquired by imitation are important ones, they may affect many aspects of individuals' lives: who they meet, how long they live, how many children they have, or whether they earn tenure. All of these factors could affect the probability that an individual becomes available as a model for others to imitate or a teacher with the opportunity to instruct the naive.

To the extent that people acquire beliefs from their parents, natural selection acts on culture in almost exactly the same way it does on genes. For example, religious beliefs affect both their survival and the reproduction of people who practice them. Sociologists Susan Janssen and Robert Hauser compared the fertility of a large sample of people living in Wisconsin.[31] Catholics (both men and women) had 20% more children, on average, than did non-Catholics. Similarly, L. McEvoy and G. Land report that members of the Reformed Latter-Day Saints Church of Missouri have age-adjusted mortalities about 20% lower than control populations belonging to other religions.[32] Behavior genetic studies indicate that religious affiliation (whether you are a Mormon or a Catholic) is culturally transmitted.[33] In Janssen and Hauser's case, people's religious beliefs are strongly correlated with the beliefs of their parents, Thus, beliefs that lead to high fertility and low mortality will increase, because people holding such beliefs are more likely to survive to adulthood and have larger families if they do, and because the children in these families will tend to have the same beliefs as their parents.

Whenever individuals are culturally influenced by teachers, peers, celebrities, and so on, natural selection acting on cultural variation can favor the increase of behaviors that increase the chance of attaining such nonparental roles. In this same scenario, when the traits that maximize success in becoming a parent are different from those that maximize success as a teacher, priest, or celebrity, natural selection acting on cultural variation can cause genetically maladaptive traits to spread.

Consider one of the most bizarre traditions in the whole ethnographic record: the existence of a subculture of people who devote more time to, and are prouder of, the length of their publication list than the number of their children. The phenomenon is potentially explicable by the effect of selection on cultural variation. We, of course, are members of this odd group and can testify to the evolutionary pressures from firsthand experience. Some of our readers will have observed university faculty at close range and may well share our experiences. To see how the selection valuing long CVs can overwhelm the complex, powerful mixture of primary and secondary urges favoring having children, consider the young assistant professor just beginning her career. Entering a new university, she needs to acquire many new beliefs or modify old ones acquired as a graduate student. She needs to know how hard to work on teaching, what the standards are by which committee work is judged, and how much time should be devoted to graduate students. And, most critical of all, how much effort should she to devote to her research? Is career advancement possible if time is also devoted to family and recreation?

In making their choices, many assistant professors decide to follow the example of older and more experienced faculty. These senior faculty represent a biased sample of the original population of assistant professors hired, because those who did not work hard and publish lots of papers were not promoted to tenure and hence aren't available to pass on their experience. Imitating tenured faculty will cause our new assistant professor to aspire to high standards in research and likely enough to postpone starting a family and limit the number of her children. This force operating on many assistant professors over several generations has produced a population that puts very high value on publications and substantially curtails childbearing. Note that we have simplified the whole story here. Throughout the educational career of our aspiring professor, she has been exposed to teachers who have faced similar career/family dilemmas, and the most successful and most influential will have been mainly those who favored career. She is liable to have fallen in love with one of her ambitious graduate school peers who shares the same background socialization and career ambitions. A successful midcareer anthropologist of our acquaintance describes the sympathetic concern of her African friends. So proud of their big families, they could not comprehend that a healthy woman would "freely" choose to have but one child.[34]

Selection for successful research faculty is driving behavior in a quite different direction from what we would predict if it were acting on genes. The role of tenured faculty member is a kind of cultural parent and social selection agent rolled into one. Potentially, natural selection on cultural variation can select for success in any role that is active in cultural transmission—biological parent, friend, leader, teacher, grandparent, and so on. The biological system is much simpler in this regard, as long as we stick with conventional organisms. There are only two roles, male and female, to worry about, and both parents make equal contributions of genes to the offspring. There are many patterns of genetic transmission that lead to the same general sorts of complexities as culture, such as Y chromosomes (transmitted from fathers to sons) and mitochondrial DNA (transmitted only by mothers),[35] but nothing quite like human culture.

Of course our young assistant professor will also take her own preferences into account as she makes decisions. If she is ambivalent about having children, she may readily adopt the publish-or-perish mentality of her most ambitious colleagues. If she is very eager to have children, she will hope that her tenure committee is more impressed by quality than quantity, and think about starting a family soon. The effect of preferences that bias decision making will lead to biased transmission. If the bias is strong, the effect of selection on the pool of models will have little effect. Plausibly, however, the bias will be weak in this case. In deciding how much time to devote to their families, young professionals must estimate not only the immediate effect on their careers and home lives but also the long-run effects on the development of their children. Biological urges to have children may be satisfied by having one or two, and the urge to achieve professional success seems to tap deeply felt biases as well. In such cases the information available to individuals may be very poor, and sentiments conflicted. Plausibly, aspiring academics will rely almost entirely on traditional beliefs, and if they do, the selective process that winnows tenured faculty will have an important effect on how faculty behave.

Why distinguish selection and biased transmission?

Biased transmission occurs because people preferentially adopt some cultural variants rather than others, while selection occurs because some cultural variants affect the lives of their bearers in ways that make those bearers more likely to be imitated. Almost every other author who has written about this topic, including biologists Luigi Cavalli-Sforza and Marcus Feldman, Richard Dawkins, and anthropologist William Durham,[36] describes biased transmission as a form of selection, often using the term *cultural selection, and* this is not unreasonable—biased transmission is a process of selective retention. Human populations are culturally variable. Some variants are more likely to be imitated than others, and thus some variants have higher relative "cultural fitness."

Nonetheless, we think that distinguishing between biased transmission and *natural* selection is very important. Biased transmission depends on what is going on in the brains of imitators, but in most forms of natural selection, the fitness of different genes depends on their effect on survival and reproduction, independent of human desires, choices, and preferences. We can understand the evolution of beak morphology in birds by asking how beaks of different size and shape affect the bird's ability to acquire food. True, we need to know something about other aspects of the bird's phenotype, so the fitness of genes affecting beak size does depend on other genes, but the dependence is much weaker than for biased transmission. Biased transmission is more like a genetic evolutionary process called meiotic drive in which "driver" genes cause the chromosomes carrying them to be disproportionately likely to be incorporated in eggs and sperm. Meiotic drive is clearly a form of selection, but most biologists think that it is useful to distinguish it from plain vanilla natural selection.

We think that the same kind of distinction should be made in the case of cultural transmission. Consider something such as acquiring an aversion to addictive drugs. If this bias is common, it will tend to suppress the spread of addiction. But even people with biases against drugs may sometimes be tempted and succumb to an addiction that could land them behind bars, or otherwise remove them from the pool of people who exercise strong cultural influence on others. Both effects may be quite important in keeping rates of drug addiction down. The aversion to addictive substances is an example of biased transmission, while the processes that influence the number of addicts available as models exemplify selection. Although distinguishing the effects of biased transmission and selection in specific empirical cases is not always easy, the distinction is important, because these processes often lead to very different evolutionary outcomes.

In our experience, most people's intuition is that psychological forces like biased transmission are much more important than natural selection in cultural evolution. They feel in control of their culture and believe they came by most of it by choice. But the truth is, we often have much less choice than we think. As Mark Twain put it,

> We know why Catholics are Catholics; why Presbyterians are Presbyterians; why Baptists are Baptists; why Mormons are Mormons; why thieves are thieves; why monarchists are monarchists; why Republicans are Republicans and Democrats, Democrats. We know that it is a matter of association and sympathy, not reasoning and examination; that hardly a man in the world has an opinion on morals, politics, and religion that he got otherwise than through his associations and sympathies.[37]

The most important hypotheses we entertain in this bookstand or fall on this issue. If the psychological forces are overwhelmingly important, then the causes of cultural evolution will ultimately trace back to innate primary values. Evolutionary "design" ultimately rests on what selection has favored and if proximal psychological forces are all that is important in cultural evolution, selection will have to do its designing on the innate aspects of psychology. Culture in this case is just a proximate cause of behavior, even if an unusual and complex proximate cause. Most evolutionary social scientists think this is the correct view. However, if natural selection sometimes acts with appreciable strength directly on cultural variation then it is also responding to the ultimate cause. Perhaps Durham's culturally transmitted secondary values are not always, strictly speaking, secondary at all. And so we will argue!

Population thinking is useful even if cultural variants aren't much like genes

Adopting a Darwinian approach to culture does not mean that you have to also believe that culture is made of miniscule, genelike particles that are faithfully replicated during cultural transmission. The evidence suggests that sometimes cultural variants *are* somewhat genelike while at other times they are decidedly not. But—and this is a big but—in either case, the Darwinian approach remains useful.

You are forgiven if you find this assertion surprising. Over the last decade or so, a lot of ink has been spilled in discussions of whether cultural variants are genelike particles. On one side of this debate are "universal Darwinists" like evolutionary biologist Richard Dawkins, philosopher Daniel Dennett, and psychologist Susan Blackmore. These authors sometimes seem to be arguing that genelike replicators are necessary for adaptive evolution, and they also think that cultural variants, which they refer to as memes, are discrete, faithfully replicating genelike particles. Because cultural variants are genelike, Darwinian theory can be applied to cultural evolution, more or less unchanged.[38] On the other side are a diverse group of critics like the anthropologists Dan Sperber and Christopher Hallpike, who argue that cultural variants are not particulate and are not faithfully replicated, so Darwinian ideas of variation and selective retention cannot be used to understand cultural evolution.

We don't agree with either side in this argument. We heartily endorse the argument that cultural evolution will proceed according to Darwinian principles, but at the same time we think that cultural evolution may be based on "units" that are quite unlike genes. We encourage you not to think of cultural variants as close analogs to genes but as different entities entirely about which we know distressingly little. They must be genelike to the extent that they carry the cultural information necessary to create cultural continuity. But, as you will see, this can be accomplished in most un-genelike ways.

The modest requirements for the properties of cultural variants are a potent rejoinder to those who believe that we can't theorize about cultural evolution until we understand exactly what cultural variants are like. If the truth is that adaptive evolution depends critically on the units of transmission, Darwin and all his followers would still be marking time, waiting for the developmental work definitively showing how genes give rise to the properties of organisms. Understanding how complexes of genes interact in development to create the traits upon which selection falls is a current hot topic in biology, if not *the* hot topic. Darwin had a very un-genelike picture of how organic inheritance worked, complete with the inheritance of acquired variation. He nonetheless did remarkably well, because the essential Darwinian processes are tolerant of how heritable variation is maintained. For the same reason, we can black-box the problem of how culture is stored in brains by using plausible models based on observable features that we do understand, and forge ahead.

Cultural variants are not replicators

In his book, *The Extended Phenotype,* Richard Dawkins eloquently argues that cumulative, adaptive evolution depends on the existence of what he calls "replicators"—entities that reproduce faithfully, that are long enough lived to affect the world, and that can increase in number. Replicators give rise to cumulative, adaptive evolution because they are targets of natural selection. Genes are replicators—they are copied with astounding accuracy, they can spread rapidly, and they persist throughout the lifetime of an organism, directing its machinery of life. Dawkins thinks that beliefs and ideas are also replicators, and coined the term *meme* to describe a cultural replicator. Memes, Dawkins thinks, can be reproduced, copied from one mind to another, thereby spreading through a population, controlling the behavior of people who hold them.[39]

We doubt that beliefs and skills are replicators, at least in the same sense that genes are. As has been forcefully argued by the cognitive anthropologist Dan Sperber,[40] ideas are not transmitted intact from one brain to another. Instead, the cultural variant in one brain generates some behavior, somebody else observes this behavior, and then (somehow) creates a cultural variant that generates more or less similar behavior. The problem is that the cultural variant in the second brain is quite likely to be different from that in the first. For any phenotypic performance there is a potentially infinite number of rules that would generate that performance. Information will be *replicated* as it is transmitted from brain to brain only if most people induce a unique rule from a given phenotypic performance. While this may often be the case, genetic, cultural, or developmental differences among people may cause them to infer different cultural variants from the same observation. Language no doubt helps get many ideas from one person to another accurately, but words are subject to multiple interpretations. As teachers, we struggle mightily to be correctly understood by our students, but in many cases we fail. To the extent that these differences shape future cultural change, the replicator model captures only part of cultural evolution.

The generativist model of phonological change illustrates the problem. According to the generativist school of linguistics, pronunciation is governed by a complex set of rules that takes as input the desired sequence of words and produces as output the sequence of sounds.[41] Generativists also believe that adults can modify their pronunciation only by adding new rules that act at the *end* of the chain of existing rules. Children, on the other hand, are not so constrained and instead induce the simplest set of grammatical rules that will account for the performances they hear. Although the children's rules produce the same performance, they can have a different structure, and therefore allow further changes by rule addition that would not have been possible under the old rules.[42]

The following example[43] illustrates how this phenomenon might work. In some English dialects, people pronounce words that begin with *wh* (*whether*) using what linguists call an "unvoiced" sound, while they pronounce words beginning with *w* using a voiced sound (*weather*). (Unvoiced sounds are produced with the glottis open, resulting in a breathy sound, while voiced sounds are produced with the glottis closed, causing a resonant tone.) People who speak these dialects must have mental representations of the two sounds and rules to assign them to appropriate words. Now suppose that people in such a population come into contact with other people who only use the voiced *w* sound. Further suppose that this second group of people is more prestigious, and people in the first group modify their speech so that they, too, use only voiced *w*'s. According to the generativists, they will accomplish this change by adding a new rule that says, "Voice all unvoiced *w*'s." So, when Larry wants to say "Whether it is better to endure . . .," the part of his brain that takes care of such things looks up the mental representations for each of the words in this sentence, including *whether* with an unvoiced *w* (because that is the way Larry learned to speak as a child). Then, after any other processing for stress or tone, the new rule changes the *w* in *whether* to a voiced *w*. In the next generation, children never hear an unvoiced *w* and adopt the same underlying representation for *whether* and *weather*. Thus, even though there is no perceptible difference in the speech of parents and children, their cultural variants differ. This difference may be important, because it will affect further changes. For example, if linguistic rules were truly replicated, future generations might recover unvoiced pronunciation of the *wh* words, whereas if they are copied from behavior, all distinctions between *wh* and *w* words will have been lost.

Replicators are not necessary for cumulative evolution

Dan Sperber and his colleagues cognitive anthropologists Pascal Boyer and Scott Atran have argued that because cultural variants do not replicate, cumulative cultural evolution is unlikely to result from the selective retention of cultural variants. They believe that the transformations that arise during cultural transmission are usually so large as to swamp the relatively weak evolutionary forces like biased transmission and natural selection.

This argument comes in two different flavors: Sometimes, Sperber and his colleagues maintain, social learning leads to systematic transformation, so that people observing a variety of different behaviors tend to infer the same underlying cultural variant. Sperber refers to such preferred variants as "attractors," because systematic transformations create a new nonselective force that moves the population toward nearby attractors. He thinks that this process is usually so strong that selective processes can be ignored.[44] In other situations, Sperber argues that the transformations that occur during social learning are unsystematic, so that people observing the same behavior infer very different cultural variants; consequently, cultural replication is so noisy and inaccurate that weak selective forces would be swamped.[45] Let's consider each of these arguments in turn.

Weak bias and selection can be important even when guided variation is strong

In many parts of the world, agricultural landowners receive a share of the crops raised on their land in lieu of rent, a practice called sharecropping. Economic theory predicts that the landowner's share will depend on the quality of the land. Owners of high-quality land should get a larger share, because they provide a more valuable input. Since land quality varies continuously, there should be all kinds of sharecrop contracts—62.3% for the landowner, 36.8% for the landowner, and so on and so on. However, typically sharecrop contracts fall into a few simple ratios. In Illinois, for example, the vast majority of contracts are of two types: 1:1 and 2:1 for the farmer.[46] Now suppose that there is a cultural variant that is the farmer's mental representation of the optimal sharecrop contract. This could take on any share between zero and one. However, further suppose that there are attractors at simple integer ratios, perhaps because such shares are easier to learn and remember. In a particular county, the optimal share might be 1.56:1. Farmers who used this contract might be more attractive as models because they make more money, and thus biased transmission would favor a 1.56:1 contract. However, guided variation would tend to increase the frequency of 1:1 contracts, and if this force were strong compared to bias, most farmers would end up believing that the 1:1 contract is best, even though they could make more money by demanding the larger share.

This example also shows that if there are multiple attractors, weak selective forces can be important even if guided variation is overwhelmingly strong. Suppose that there are two equally strong attractors for sharecrop contracts, 1:1 and 2:1. and that a population of farmers starts out with a range of contracts. After a short while, everybody will think one of the two simple ratios is the best contract—some 1:1 and others 2:1. Because these are strong attractors, they will be transmitted extremely faithfully. People who observe somebody using a 1:1 contract will correctly infer that that person thinks even shares are the best contract. Similarly, people observing a 2:1 contract in action will correctly infer the underlying belief. If, the 2:1 contract is a little more profitable for landlords 2:1 contracts will gradually replace the 1:1 contract, because other landlords are more likely to imitate the successful. In effect, *multiple* strong attractors lead to discrete, genelike cultural variants. Only if one attractor is stronger than the sum of all the other forces acting on other attractors will guided variation completely determine the evolutionary outcome.

Adaptive evolution can occur even when transmission is very noisy

When cultural transmission is noisy, it cannot produce cultural inertia for exactly the same reasons that genetic transmission does. To see this, suppose there are only two cultural variants in some domain, labeled *A* and *B*. Each generates different but overlapping distributions of observable behavior. When cultural learning occurs, naive individuals, perhaps children, observe a sample of individuals from these distributions, make inferences, and then adopt their own mental representation. This process is very sloppy—a naive individual who observes an A infers that the individual is an A 80% of the time and a B 20% of the time. Similarly, a naive individual who observes a B infers B 80% of the time and A 20% of the time. It is clear that this kind of social learning will not lead to replication at the population level. Suppose that 100% of the people initially have cultural variant A. After one generation 80% will be A, after two generations it will be 68%, and by generation 5 or so, the population will have converged to a random distribution of cultural variants. Only very strong selection or bias could generate cumulative adaptation.

However, just because cultural transmission is inaccurate, it does not necessarily follow that there can be no cultural inertia or cumulative evolution of adaptations. Transmission processes can lead to accurate replication at the level of the population, even when individual social learning is loaded with errors. As before, suppose that every naive individual observes the behavior of a number of models and makes inferences about the beliefs that gave rise to each person's behavior, and that people make the wrong inference 20% of the time. Now, suppose that individuals adopt the cultural variant that they believe is *most common* among their models. This is a form of biased transmission, because some variants are more likely to be adopted than others. However, unlike the biases discussed above, the nature of the bias is independent of content. It depends only on which variant is more common, and represents a "conformist" bias in social learning. In the next chapter you will see that there is good evidence that people do have a conformist bias, and that there are good evolutionary reasons why this should be the case. A conformist bias at the individual level leads to reasonably accurate replication at the population level, even when individual inference about underlying mental representations is inaccurate. For example, if everyone is A, 20% of the As are mistaken for Bs, but the chances are that most naive individuals will observe samples in which A are most common as long as these samples are large. Conformist bias corrects for the effect of errors because it increases the chance that individuals will acquire the more common of the two variants.

Yet the combination of high error rates and a conformist bias does not result in the same kind of "frictionless" adaptation as genetic replication. Highly accurate, unbiased genetic replication allows minute selective forces to generate and preserve adaptations over millions of years. Error-prone cultural replication, even when corrected by a conformist bias, imposes modest, but still significant forces on the cultural composition of the population. This means that only selective forces of similar magnitude will lead to cumulative adaptation. We do not think this is a problem: the forces of bias and natural selection acting on cultural variation are probably much stronger than those that shape genetic variation because they work on shorter timescales, and are often driven by psychological processes, not demographic events. The empirical record supports this somewhat, providing examples of innovations that spread over decades, not millennia.

Cultural replication can be quite accurate

Cultural transmission does not *have* to be biased and inaccurate. In fact, sometimes arbitrary cultural variants are transmitted with considerable fidelity. Take word learning, for example. The average high school graduate has mastered about sixty thousand words—an astounding feat. Learning words is a difficult inferential problem for the reasons already mentioned. The child on the nursery floor hears the word *ball* and surveys the scene. Perhaps the adult is referring to the red ball rolling across the floor, but many other inferences are possible. It could be that the adult is referring to moving red objects, the fact that it is warm, or the fact that the ball is rolling north. Despite seemingly endless opportunities for confusion, children acquire about ten new associations between a range of sounds and a meaning every day.

According to developmental linguist Paul Bloom, children use a variety of strategies to acquire their immense vocabularies.[47] They behave as if they start with the assumption that words refer to objects, and even very young children have innate presumptions about what objects are. Our hypothetical child will interpret the red ball as an object because it is connected, bounded, and moves as a unit unless some further evidence proves otherwise.[48] "Joint attention" provides another important mechanism for learning language.[49] Children follow the gaze of adults who can often be induced to pay attention to what a child is paying attention to. In the course of these games, the adult often names the object of joint attention, usually as a part of a more complex utterance: "A red ball! I'll roll you the red ball!" To extract *ball* and *red* out of such a language stream as names of a certain kind of round object and a color that applies to many objects is quite a feat, but the potential ambiguity is sharply limited by the assumption that the utterance is only relevant to the object of joint attention, the red ball. Another strategy children use is what psychologists call "fast mapping." Suppose a three-year-old is presented with two balls, one red and one turquoise. An experimenter asks, "Toss me the chromium ball, not the red one, the chromium one!" The child knows the color term *red* very well but not *chromium* or *turquoise*. Typically the child simply assumes that *chromium* means "turquoise" and many retain this false hypothesis for at least a week. In many cases, further experience confirms hypotheses formed by fast mapping and they go on to become a durable part of the vocabulary. Grammatical cues also play a role in language learning. For example, the child knows that *red ball* is not an action from its role in the sentence. These are only a few of the mechanisms that allow kids to accurately acquire a huge vocabulary without any innate predispositions about what words mean.

Historical linguistics suggests that these mechanisms can maintain detectable similarities in languages over hundreds of generations. Sir William Jones, the Chief Justice of India, launched the discipline of historical linguistics at the end of the eighteenth century by demonstrating that Sanskrit has certain remarkable resemblances to European languages such as Greek and Latin, resemblances too numerous to be explained by chance. Instead, these languages and a variety of others belonging to the Indo-European language family are all descendants of a single language, known as Proto-Indo-European. As the people speaking this language spread out across Eurasia linguistic communities became isolated and the languages gradually diverged. Exactly how long ago this occurred is controversial. Some think that the speakers of Proto-Indo-

European were the earliest farmers who dispersed from their agrarian homeland in southwestern Asia beginning about ten thousand years ago. Others think that they were horse-mounted nomadic herders who emerged from Central Asian or southeast European regions about six thousand years ago.[50] To be conservative, let's suppose that the Proto-Indo-European was spoken six thousand years ago, or roughly 240 human generations in the past. Contemporary Indo-European languages are connected to the speakers of Proto-Indo-European by a chain of cultural transmission 240 generations long. Each generation, children learned the sound-meaning associations from adults, and then served as models for the next generation. Thus the similarities that historical linguists use to link these languages have survived 480 generations of cultural transmission, indicating that cultural transmission can be quite accurate indeed.

Cultural variants need not be particulate

Many people believe that cultural inheritance must be particulate if it is to undergo Darwinian evolution because, the story goes, only particulate inheritance conserves the variation necessary for the action of natural selection. Biology textbooks often illustrate this idea by explaining how the discovery of Mendelian genetics rescued Darwin from the problem posed by a British engineer named Fleeming Jenkin. Jenkin was nobody's fool—a longtime associate of the great but antievolutionist physicist Lord Kelvin, he played a key role in the design and construction of the first transatlantic cable and made important contributions to economics, including inventing the supply and demand curve. Nowadays, however, he is mainly known for pointing out that if inheritance works by taking the average of the parental genetic contributions, as Darwin proposed, then the amount of variation would be reduced by half each generation. Therefore, the variation necessary for natural selection to be effective would rapidly disappear. This critique vexed Darwin greatly, but it wasn't resolved until geneticists like R. A. Fisher showed that variation persists because genes don't mix; each parent's genes remain separate particles in offspring.

This story is true but misleading. Because mutation rates are very low, the particulate nature of genetic inheritance is crucial for maintaining the genetic variation. However, perhaps the analog of mutation in cultural transmission is not so low.[51] We can even imagine that cultural transmission is sufficiently noisy and error prone that blending inheritance would be an *advantage* in keeping cultural variation from growing disastrously large. In a noisy world, taking the average of many models may be necessary to uncover a reasonable approximation of the true value of a particular trait. For example, when you speak, the sounds that come out of your mouth depend on the geometry of your vocal tract. For example, the consonant *p* in *spit* is created by momentarily bringing your lips together with the glottis open. Narrowing the glottis converts this consonant to *b,* as in *bib*. Leaving the glottis open and slightly opening the lips produces *pf,* as in the German word *apfel*. Linguists have shown that even within a single speech community, individuals vary in the exact geometry of the vocal tract used to produce any given word. Thus, quite plausibly, individuals vary in the culturally acquired rule about how to arrange the inside of the mouth when they are saying any particular word. Languages vary in the sounds used, and this variation can be very long lived. For example, in dialects spoken in the northwest of Germany, *p* is substituted for *pf* in *apfel* and many similar words. This difference arose about AD 500 and has persisted ever since.[52]

Now suppose that children are exposed to the speech of a number of adults who vary in the way that they pronounce *pf*. Each child unconsciously computes the average of all the pronunciations that she hears and adopts the tongue position that produces approximately the average. There is no doubt that this act of averaging would tend to decrease the amount of variation in the population each generation. However, phenotypic performances also will vary as a result of age, social context, vocal tract anatomy, and so on. Moreover, learners will often misperceive a performance. These sorts of errors in transmission will keep pumping variation into a population as blending bleeds it away. Further note that the errors one makes will affect one's performance and will thus affect what learners use as the basis for constructing their own way of saying *pf*. Some variation will always remain if any heritable errors occur in the cultural transmission process, as surely they do.

With this sort of averaging mechanism, mental rules are not particulate, nor do they replicate. A child may well adopt a rule that is unlike any of the rules in the brains of its models. The phonological system can nonetheless evolve in a quite Darwinian way. More-attractive forms of pronunciation can increase if they have a disproportionate effect on the average. Rules affecting different aspects of pronunciation can recombine and thus lead to the cumulative evolution of complex phonological rules. In fact, this model faithfully mimics all the usual properties of ordinary genetic evolution. We are confident of this claim, because models exactly like it have been used in population genetics to represent the evolution of characters such as height that are affected by many genes, each with a small effect. They provide a good approximation to genetically more realistic models and are much easier to analyze.[53]

Cultural variants need not be small, independent bits

Many people believe that a Darwinian approach to cultural evolution requires breaking culture into little, independent bits, an anathema to many anthropologists who believe that cultures are tightly integrated systems of shared meanings. Just as the syntax of a language is made up of a system of interdependent rules, so are the cultural meanings embedded in systems of kinship, cosmology, law, and ritual. Since Darwinian models require that cultures be decomposed into independent, atomistic traits, the argument runs, Darwinian models must be wrong. For example, Christopher Hallpike complains:

> The absence of any . . . structural concepts inevitably reduces the examples of memes and culturgens to ridiculous laundry lists of odds and ends—Dawkins's tunes, catch-phrases and ways of making pots, and Lumsden and Wilson's food items, colour classifications, 6000 attributes of camels among the Arabs, and the ten-second-slow-downs by which drivers cause traffic jams.
>
> …In fact, such theories of basic units of culture do not rest on any evidence, or any sociological theory at all, but are simply proposed because if one is trying to explain culture on the basis of a neo-Darwinian theory of natural selection, it is highly inconvenient *not* to have a "unit" like a meme or culturgen, quantifications of which can be treated as continuously variable over time like the gene.[54]

This criticism misses the mark. Perhaps we (and others of our persuasion) have fostered this view by choosing very simple examples to illustrate our ideas, but there is absolutely nothing in the theory that requires that cultural variants be little bits of culture. People may choose between great, linked cultural complexes—between speaking Spanish or Guarani, or between remaining a Catholic or becoming a Seventh-Day Adventist; or they may choose between smaller, more loosely linked items of knowledge—between pronouncing *r* at the end of a word or not, or between different views about the morality of contraception. At a *formal* level, Darwinian methods will apply equally well in either case. We keep track of the different variants, independent little bits or big complexes as the case may be, present in a population, and try to understand what processes cause some variants to increase and others to decline. The same logic applies whether the variants are individual phonological rules or entire grammars.

Cultures are not tightly structured wholes

Whether cultures actually *are* tightly integrated wholes is an important empirical question. While there has been surprisingly little systematic attention paid to this problem, a great mass of observational data bear on it. We believe that these data suggest that culture is a complex mixture of structures. Some cultural variants are linked into coherent wholes, while others float promiscuously from culture to culture.

The data from linguistics suggest that even the tightly interlinked rules underlying language sometimes diffuse and recombine. Words, phonological rules, and syntax all can diffuse and recombine independently, and as a result, different components of a single language often have a different evolutionary history. You can see this in the history of English. Some words in the English lexicon are derived from French, while others come from German. In German, the object sometimes comes before the verb in a sentence, but in French the object always follows the verb. English adopts the French syntax, although the majority of spoken English vocabulary is derived from German. Most English phonology is descended from a Germanic language; but unlike German speakers, English speakers distinguish [v], as in *veal,* from [f], as in *feel,* apparently as a result of the influence of Norman "loan words." linguists Sarah Thomason and Terrence Kaufman[55] provide many examples from other languages, including the Ma'a language spoken in northern Tanzania that has a basic lexicon related to Cushitic languages and a grammar related to Bantu languages. They summarize by saying that "any linguistic feature can be transferred from any language to any other language."[56] They go on to argue that it is the actual pattern of social, political, and cultural interaction that determines the extent and kinds of diffusion among languages.

While the linguistic data suggest that any linguistic feature can diffuse from one language to another, they also suggest that the rate at which different features diffuse depends on a number of linguistic and social factors. What linguists call "typological distance" seems to be the most important linguistic factor. Typological distance measures the extent to which two languages have similar structure. All other things being equal, the more similar two languages are the higher the rate of borrowing. In turn, more highly structured subsystems of language diffuse and recombine at a slower rate than less structured systems. Individual words are more or less independent of each other, and as a result, they are the first items to diffuse when two languages come into contact. Inflectional morphology (for example, different verb forms that depend on the person, timing, or type of action) is linked in a complex, multidimensional system and therefore will diffuse very slowly unless the inflectional morphology of neighboring languages shares a similar structure.[57] For example, Norse had a substantial impact on English grammar even though only a small number of Danes occupied a small part of England for a relatively short time, because the typological distance between Norse and Old English was small. The rate and direction of diffusion is also strongly influenced by many social factors, the extent of bilingualism, the context in which bilingual speakers use each language, and the relative prestige of groups speaking different languages.[58]

Good evidence also suggests that language is not a good predictor of material culture—anthropological jargon for the kinds of tools, containers, dwellings, and clothing that people use. One recent study compared the artifacts collected at a number of villages on the northern coast of New Guinea during the early 1900s with the languages now spoken in those villages.[59] There was no association between language spoken and the kinds of artifacts used when the distance between villages was held constant. This means that the material cultures of two villages thirty kilometers apart with closely related languages are no more similar than the material culture of two villages thirty kilometers apart in which completely unrelated languages are spoken. Studies in Africa and North America come to the same general conclusion.[60]

A vast amount of anecdotal data provides circumstantial evidence that other components of cultures are a mix of loosely and more tightly linked elements. There are obviously many examples of important cultural similarities and differences that do not map onto linguistic differences. For example, male and female genital mutilation are common customs throughout central and East Africa and are practiced by people who speak very distantly related languages. California acorn-salmon hunter-gatherers and corn farmers of the Southwest both encompassed diverse language groups. The spread of religious practices, including the spread of the Sun Dance on the Great Plains, Islam across central Asia, and millenarian movements in Melanesia, along with the contemporary spread of Protestantism in Latin America, provide additional examples of cultural practices diffusing across many different cultures/languages. On the other hand, that ritual practices and systems of religious belief can be identified as they diffuse among widely different cultures suggests that the many beliefs that make them up *are* reasonably tightly integrated and as a result *do* cohere. Some scholars, such as philologist Georges Dumézil,[61] argue that cultures have a set of core beliefs, and these core beliefs create cultural continuity over thousands of years.

Population thinking helps explain variation in cultural coherence

That cultures are not made up of independently evolving bits but are composed of at least partly integrated complexes of beliefs and values is not an embarrassment for the Darwinian approach. Quite to the contrary, population-based evolutionary theory has tools to help us think clearly about the degree, pattern, and process of integration. What we mean by integration here is that the various components of a particular aspect of culture covary in space or time for particular reasons. Because a population-based theory of culture focuses on patterns of variation, it also provides a natural framework to describe patterns of integration.

Sometimes the existence of one variant doesn't create any bias for or against other variants. Such is often the case for lexicon. You can use the Spanish loanword *arroyo* for a dry gully without also having to adopt *gato* for cat. In this case, the mixing of individuals from different populations has a powerful tendency to erase differences between populations, destroying any structure that previously existed. On the other hand, the effect of mixing is limited if you learn one set of things from one person and other sets of things from others. This may produce independent subcultures within a population, subcultures that can even coexist within a single individual. For example, the subculture of science is reasonably coherent and coexists with the subculture of rock climbers, and in English-speaking countries both groups share the same language. There are even a few scientists who climb rocks and speak English, but they certainly don't form a subculture of rock-climbing, English-speaking scientists—especially if scientists who climb rocks make no special effort to recruit their students to become rock climbers or to persuade their rock-climbing buddies to become scientists. Being a scientist may have no impact on your success as a rock climber and no more impact on your social status than having any one of a number of other elite occupations. In this case, evolutionary processes will have independent effects on each of the three trait complexes. The evolution of some traits can be substantially decoupled from the evolution of others.

When the interaction between elements is strong, biased transmission can build coherence even in the face of substantial mixing pressure. Suppose that rock climbing has the effect of enhancing cognitive skills that are particularly useful for physical environmental scientists (geologists, meteorologists, and the like), although they detract from one's ability to be a good biologist and teach exactly the wrong lessons for social scientists. Rock climbers would then tend to be especially successful environmental scientists, but very poor social scientists. If successful environmental scientists tend to attract more students who learn both science and rock climbing from their mentors, a correlation between rock climbing and environmental science would arise. On the other hand, few successful social scientists would be rock climbers, and wouldn't encourage this hobby among their students. Successful social scientists might be prone to, say, play soccer. Eventually, a complex of coherent traits may arise which separates physical and social scientists. The gulf between the physical and social sciences is real, although we have no reason to think that rock climbing or soccer played any role in their estrangement!

Why bother with evolutionary models?

Evolutionary models aren't the only way to study how human behavior and human societies change through time. Historians, and historically minded scholars in other disciplines, have long studied social change without any reference to evolution, evolutionary forces, or anything of the like. Instead, historians seek to generate a reliable narrative account of particular sequences of historical events, and have developed rigorous methods for answering questions like What motivations led the Continental Congress to declare American independence in 1776?[62] The goal is a true historical narrative of events. Historians typically eschew simple abstract models that can be applied to a variety of cases. Instead, they focus their efforts on developing a rich explanation of events within a particular historical frame. This approach is without doubt successful in accounting for temporal change in human societies, so a reasonable question is, why should we abandon it in favor of simple, process-based models?

We think the question so phrased is misleading. You don't have to choose between simple abstract models and rich historical explanation—these modes of explanation are complementary not competing.[63] Historians are certainly right: every concrete problem in cultural evolution is embedded in a complex, historically contingent frame, and all causes of events are local to that frame. However, the same is true for genetic evolution—the evolutionary biologist knows complexity and diversity as intimately as the historian. Biologists are responsible for millions of species with a huge range of characteristics and complex histories, and for the interactions of many species in complex communities. Successful field biologists typically have steeped themselves in natural history from their teenage years onward.[64] If they followed the practice of many historians and anthropologists, they would give up the concept of natural selection and speak simply in terms of the concrete events in the lives of particular organisms living in particular places and particular times that caused some genes to spread and others to diminish. After all, these local causes are all that natural selection can ever amount to in concrete terms.[65]

Instead, these very same biologists typically have a love of simple explanatory models. What gives? The answer is that such explanatory models are not laws but tools to be taken up or not as the situation warrants. Good models are like good tools: they are known to do a certain job reasonably well. Simple models that work well for a wide variety of jobs are an especially valuable part of the biologist's tool kit.

Having a toolbox filled with such models brings three important benefits. First, it is economical. The complexity of any interesting problem is likely to demand more hard thinking than any given investigator can bring to bear by himself. Person-months, if not person-years, have gone into the development of existing models, and no single investigator is likely to develop anything half as good on the spot. A mechanic who insisted on building all his tools from scratch could not be nearly as productive as one who shops at the hardware store. When available models don't work, the reasons they don't provide clues about what to try next, usually a modification of an existing model.

Second, simple models provide islands of conceptual clarity in the midst of otherwise mind-numbing complexity and diversity. Although this is not a book about formal models of cultural evolution,[66] our thinking about the major issues in cultural evolution is schooled by mathematical formalism borrowed from population genetics, game theory, and economics. These three disciplines share an enthusiasm for *simple, general* models. And these models can prevent serious errors in reasoning—errors that are all too frequent in disciplines that eschew such models.[67]

Third, by using a standardized conceptual tool kit, we increase the chance that we will detect useful generalizations in spite of the complexity and diversity of human behavior. Evolutionary biology and ecology are not without encouraging results in this regard. Although historical contingency and local uniqueness clearly matter, we can detect some general patterns in the worlds we study.[68] From the theory-as-tool kit perspective, every study provides a bit of information about the circumstances in which specific tools succeed or fail. Your colleagues provide the tools to carry to the work, and you in turn provide what help you can to the investigator with a similar problem by explaining which tools worked for you and which did not. Science advances by developing better methods, and an expanding set of empirically useful theoretical models.[69]

Darwinian tools help get the right answer

We are advocating that social scientists change the way they do business, *supplementing* their usual tool kit with ideas imported from biology. Naturally enough, many of them resent unsolicited advice from outside their disciplines. The philosopher Elliot Sober has captured one common reaction in a paper in which he argues that population-based models of cultural change will be of little interest to social scientists, because cultural evolution depends on learning rules.[70] As he puts it,

> My main reason for skepticism is that these models concern themselves with the *consequences* of transmission systems and fitness differences, not with their *sources* [his emphasis].[71]

To understand why some ideas spread but others do not, you need to know people's learning rules, their transmission biases, and the like. Why did someone invent a given cultural variant in the first place? Why is it attractive to others? You have to know which ideas will be imitated and which will be ignored. This knowledge does not come from within the Darwinian model, Sober argues; rather, it has to come from some other theory. Given learning rules, Darwinian models can predict the trajectory of cultural change, but according to Sober, this is of much less interest to social scientists than people's preferences. In other words, Sober thinks that population-based theories take all the important stuff as given, and concentrate on the stuff that nobody really cares about. The hard parts of social science don't involve its population-level properties, and the population level, unlike the biological case, is trivial. This critique has in common with many others the idea that cultural evolution is somehow so different from organic evolution that population-level processes simply don't matter.

There are three things wrong with this argument. First, it assumes that content-driven biases are the only important process affecting cultural change, and this is simply false. Biases are important, but so are processes like natural selection, which can only be understood in terms of the population dynamics of alternative cultural variants. Second, it assumes that once you know people's learning rules, how they make choices about which culture to imitate and perform, it's easy predict the evolutionary outcome. Or, in other words, we are all good intuitive population thinkers. Much experience in the relatively simpler world of evolutionary biology suggests that this is not the case. Finally, the biases are themselves the result of interacting genetic and cultural evolutionary processes Understanding the evolution of the rules requires a theory that can work out how rules influence the social environment, which in turn influences what social information is available.

Conclusion: We are ready to get to work

We have now introduced you to all of the essential components of the Darwinian analysis of cultural evolution.

The basic steps of Darwinian analysis are

- draw up a model of the life history of individuals;
- fit an individual-level model of the cultural (and genetic, if relevant) transmission processes to the life history;
- decide which cultural (and genetic) variants to consider;
- fit an individual-level model of the ecological effects to the life history and to the variants;
- scale up by embedding the individual-level processes in a population; and
- extend over time by iterating the one-generation model generation after generation.</bl>

In a theoretical model, the final product will contain mathematical terms and operations representing each of these steps. For a large set of models built on these principles, see our earlier book and works in the same genre.[72] In an empirical investigation, we want descriptions and measurements of as many of these components as we can manage.

In order to actually make progress with theoretical or empirical work, you have to be willing to simplify, simplify, and then simplify some more. The Darwinian tradition encourages us to modularize problems and deal with highly simplified bits of nature one at a time. We are fond of simple models that are deliberate caricatures of the real world. We are also fond of abstract experiments that admit only a tiny a bit of realism of the real world. We are fond of field data that clearly show the effects of one process and hate data where several processes interact to produce an unintelligible mishmash. We don't have these preferences because we think that the real world normally resembles these kinds of simple models, experiments, and field situations. No sensible scientist thinks that the complexity of the organic or cultural world can be subsumed under a few fundamental laws of nature or captured in a small range of experiments. The "reductionism" of evolutionary science is purely tactical. We do what we can do in the face of an awesome amount of diversity and complexity. Simple, deliberately unrealistic models and highly controlled experiments have great heuristic value, because they capture manageable bits of realism. We use them to school our intuitions. We undertake empirical studies looking at limited aspects of a phenomenon—technology, politics, or art, say—because we haven't the mental or physical resources to do more. We look for the simplest real cases we can find to develop some confidence that our models and experiments are at least sometimes true.[73]

We hope your mind is racing ahead, anticipating the modifications and extensions to this rudimentary map of cultural evolution. If so, you may well already be in uncharted territory. The possible avenues of exploration are large relative to those traveled thus far. In what follows, we will repeat the exercise of this chapter for several more forces of cultural evolution, examine the results of the models in light of the current evidence, and sketch what we believe is a basic picture of the cultural evolutionary process in humans. If we don't do justice to your favorite regions, we aim to leave you with the tools for doing so at home. You can't hurt yourself.

---

1. Attributed to Charles Anderson Dana (1819–97), longtime editor of the New York *Sun* and one of the creators of the modern American newspaper. Dana is the source of another quotation which we admire: "Fight for your opinions, but do not believe that they contain the whole truth, or the only truth."
2. Burrow (1966) provides a classic account. See also Richerson and Boyd 2001a.
3. White 1949; Sahlins, Harding, and Service 1960.
4. Steward 1955, Sahlins Harding and Service 1960, Harris 1979.
5. Johnson and Earle 2000; Carneiro 2003.
6. See Sahlins, Harding, and Service 1960 and Steward 1955 for two approaches to dealing simultaneously with the complexity of the evolution of particular traditions and the general trend. For an authoritative modern treatment of this kind of evolutionism, see Johnson and Earle 2000.
7. E.g., Cohen 1977 on the origin of agriculture. Harris 1977, 1979 and Johnson and Earle 2000 make population pressure the engine of cultural evolution.
[8] Richerson, Boyd, and Bettinger 2001, Richerson and Boyd 2001c.
9. Blurton-Jones and Konner 1976.
10. Merlin Donald 1991 takes the invention of literacy and similar information technologies to be one of the major revolutions in the origins of the modern mind, greatly increasing the accuracy and volume of information that individuals can access. We don't mean to discount information technology! The pioneering evolutionary economists Richard Nelson and Sidney Winter 1982 used firms as their unit of analysis and the routines of firms as the unit of culture. In the fifth chapter of their book they give the best discussion we know of on the means by which culture can be carried outside individuals' heads.
11. See Griffiths 1997 and Wierzbicka 1992 for the case that the scientific study of emotions has been handicapped by culture-specific concepts. Richard Nisbett 2003 presents considerable evidence that Asians think quite differently from Americans.
12. See Baum 1994 for an evolutionarily sophisticated version of behaviorism and Pinker 1997 for the cognitivist approach.
13. Gallistel 1990 on mental representations, Churchland 1989 on why not.
14. Jackendoff, 1990 commentary on Pinker and Bloom 1990.
15. To be more precise, social learning is itself a concept with several subconcepts, only some of which would support imitation-based culture in the human sense (not to say that humans don't

sometime use simpler sorts of social learning). See Galef 1988 for an introduction to these complexities.

16. The use of such toy models for didactic purposes is a common practice in some disciplines (e.g., economics, evolutionary biology) but not others (e.g., anthropology, history).

17. Atran 2001; Boyer 1998; Sperber 1996.

18. Salamon 1992, 172.

19. The intent of the distinction between transmission and forces is analytical, not ontological. It is often convenient to assume that there is one step in the life history in which perfect transmission occurs, followed by another step in which the mind applies biases to select among perfectly learned cultural variants. The staged life history is a trick to simplify the structure and analysis of models of evolution that are borrowed from evolutionary biology. The facts may be quite different; the bias may be applied at the point of learning to distort the cultural variant as it is learned. Under most conditions, slight differences in structure don't affect the outcome of models, so we claim that the step-structure approach is usually an innocent simplification. In theory, and no doubt occasionally in practice, there will be cases where a more realistic psychology of transmission is absolutely necessary. A failure to distinguish between tactical analytical simplifications and truth claims has led a number of unwary critics of dual inheritance theory to unwarranted conclusions. For example, it may seem "reductionistic" to analyze the no doubt very complex events occurring in farming communities in the Midwest in terms of two cultural variants and two forces. We claim no more than that the highly simplified picture we present is a tolerably good first approximation to that complex phenomenon. Additional variants and forces would be necessary to explain even Salamon's data, much less all the facts of the case, assuming (counterfactually) that they could all be put on the table. In truth, no empirical or theoretical study can manage more than a modest fraction of all the processes ongoing in particular cases of evolution. One is stuck with a choice among alternative simple models (or simple experimental designs) and between doing analysis and practicing mysticism. At least in favorable cases, a few things do dominate the evolutionary process and our analysis leads to great insights. We hasten to add we imply no objection to mysticism. Many excellent "hard" scientists become mystics after two beers; Darwin's last paragraph of the *Origin* is a first-class example. In unfavorable cases there is little to do but be in awe of the complexity of the tangled bank. In Richerson and Boyd 1987 we outline and defend the simple-models-of-complex-phenomena strategy employed by evolutionary biologists, economists, and engineers, among many others.

20. For a technical discussion see Boyd and Richerson 1985, chap. 5.

21. Ryan and Gross 1943. Rogers 1983 surveys this literature, counting 3,085 studies from 10 different disciplines as of that date.

22. Rogers with Shoemaker 1971 showed that perceived advantage was one of the commonest effects in studies of the diffusion of innovations. This book did a primitive meta-analysis of some fifteen hundred diffusion-of-innovation studies. Henrich 2001 shows how a quantitative analysis of such adoption data can be used to estimate the influence of the various forces of evolution.

23. Wiessner and Tumuu 1998; Yen 1974. See Crosby 1972, 1986 for a discussion of the rapid spread of many New World crops in the Old World following the voyages of Columbus, and of Old World plants and animals in the New.

24. Labov 1994 discusses the principles internal to the structure of language that help drive linguistic evolution.
25. Durham 1991.
26. Lindblom 1986, 1996.
27. Alexander 1979; Lumsden and Wilson 1981.
28. As Melvin Thorpe, the fictional governor of Texas in *The Best Little Whorehouse in Texas,* puts it: "Ooo . . . I love to dance the little sidestep. Now they see me now they don't, I've come and gone. And . . . Ooo I love to sweep a-round a wide step, cut a lit-tle swath and lead the peo-ple on" (from "The Side Step," lyrics by Carol Hall). Labov 1994 describes many cases in which language change due to psychological factors decreases communication efficiency.
29. The idea that ideas compete and that the results of this competition drive human history was elaborated by sociologist Gabriel Tarde 1903 at the turn of the twentieth century.
30. Castro and Toro 1998 discuss the potential importance of teaching as opposed to simple imitation in the evolution of some important features of human cultures.
31. Janssen and Hauser 1981.
32. McEvoy and Land 1981.
33. Eaves, Martin, and Eysenck 1989.
34. Some social scientists propose that an explanation of cases like this in terms of beliefs, desires, and intentions is sufficient. We disagree. The explanation that assistant professors work hard because they intend to get tenure and tenured faculty vote against those who are slack because they intend to maintain the quality of the department has great intuitive appeal to our folk psychology. How, then, do we explain why professors tend to prefer writing papers to having children, while rural Africans have quite a different set of intentions? At best,beliefs, desires and intentions are proximal explanations themselves in need of an ultimate evolutionary explanation. See Rosenberg 1988 for a critique of scientific explanations in terms of folk psychology.
35. Hamilton 1967; Dawkins 1982; Jablonka and Lamb 1995; Rice 1996.
36. Cavalli-Sforza and Feldman 1981; Dawkins 1976; Durham 1991.
37. "Cornpone Opinions," Twain 1962, 24.
38. See Blackmore 1999 for a review of the work done using the meme concept. Richard Dawkins's foreword to Blakemore's book gives a particularly clear example of how important the high fidelity of transmission is taken to be by Dawkins at least. See Durham and Weingart 1997 for a discussion of alternative proposals for the unit of cultural inheritance. Dennett 1995 in *Darwin's Dangerous Idea* , provides an extended argument in favor of the idea that replicators are necessary for cumulative adaptation.
38. See Aunger 2002 for an elaboration and critique of this view.
39. See ibid. for an elaboration and critique of this view.
40. Sperber 1996.
41. Bynon 1977, characterizing scholars like Chomsky and Halle 1968.
42. Note that this phenomenon may take some of the bite out of Chomsky's argument from the poverty of the stimulus. Perhaps in the case of grammar, all native American English speakers don't all have the same rules in their heads. Perhaps learners adopt the first rule that they stumble across that generates grammatical sentences an acceptably large percentage of the time. There may be more than one rule that does so, so no one is really speaking exactly the same language.

Individual speakers certainly do have small differences in their speech called ideolects. According to sociolinguists, ideolectual variation is the raw material out of which language evolution grows, a quite Darwinian notion (Labov 2001; Wardhaugh 1992). It is not so clear whether ideolect includes grammatical rules, but if it does, the sociolinguist's picture of the evolution of phonology may extend to syntax.
43. Bynon 1977.
44. Sperber 1996: Chapter 5.
45. Sperber 1996; Boyer 1998, 1994; Atran 2001.
46. Burke and Young 2001. In addition to the 1:1 and 2:1 contracts, they also observed a small number of 3:2 contracts, and, even among the highly market-oriented farmers of Illinois, virtually no other shares. Burke and Young also show that farmers don't adjust shares by varying other inputs such as fertilizer or pesticides.
47. Bloom 2001.
48. Spelke 1994.
49. Tomasello 1999.
50. Mallory 1989.
51. Lande 1976.
52. Bynon 1977.
53. Cavalli-Sforza and Feldman, 1976, 1981, Karlin 1979**.**
54. Hallpike 1986, 46.
55. Thomason and Kaufman 1988. See also Thomason 2001.
56. Thomason and Kaufman 1988.
57. Ibid.
58. Ibid.
59. Welsch, Terrell, and Nadolski 1992.
60. Jorgensen 1980; Hodder 1978.
61. Dumézil 1958; Hallpike 1986; Mallory 1989, chap. 5.
62. Brown 1988. Vayda 1995 argues that such explanations are much to be preferred to the general process accounts that we shall focus upon.
63. Boyd and Richerson 1992a.
64. Darwin "wasted" his college years following his dogs across the countryside, shooting birds, collecting beetles, and speculating about geology under the guidance of Adam Sedgwick. The contextual detail such a naturalist commands certainly rivals that of ethnographers and historians. Some naturalists write and speak lyrically about the pleasure they get from looking the complexity and diversity of nature in the eye. E. O. Wilson's 1984 celebration of the naturalist's craft, *Biophilia* 1984, is an excellent example. So is Darwin's last paragraph of the *Origin* and many passages in his *Journal of Researches* (*Voyage of the* Beagle). W. D. Hamilton was before all else an intrepid and perpetually entranced naturalist, according to those who knew him best. The same can be said of the dean of living evolutionary theorists, John Maynard Smith. One of us (Richerson) has spent quite a lot of effort trying to understand the ecology of lakes, one of the simplest sorts of ecological systems, and will trade stories of complexity and diversity with any human scientist who cares to defend the idea that our species is any more complex than the average ecosystem. An accessible description of how evolutionary biologists immerse themselves in the detail of their chosen "system" is Jonathan Weiner's 1994 book *The Beak of*

the Finch , describing Peter and Rosemary Grant's wonderful study of the evolution of Darwin's finches on the Galapagos Islands. This is a high-end study, to be sure, but every serious field study of evolution at least aspires to something like its resolution of the concrete events that are eventually summarized as selection of a certain strength on a certain trait.

65. Other authors, using vague or different arguments, have tried to make the case that something different about cultural and genetic evolution has led to cultural evolution being properly studied with methods quite different from genetic evolution. We think that, because of the general similarity of the evolution of the two systems, they all fall prey to this "sauce for the goose, sauce for the gander" analysis. One outfit or the other is doing something wrong! See Sober 1991 and Marks and Staski 1988.

66. See Boyd and Richerson 1985 and Cavalli-Sforza and Feldman 1981 for a full mathematical treatment of the issues. We refer extensively to these and other formal theoretical studies in later chapters.

67. For example, archaeologists often use population pressure to explain phenomena on very long timescales, such as the origins of agriculture. A little elementary modeling of demographic and evolutionary timescales suggests that the changes in subsistence that led up to and then resulted in domestication of plants and animals happen so *slowly* that demographic processes cannot explain either their occurrence in time nor their rate of change (Richerson, Boyd, and Bettinger 2001). As in organic evolution, population pressure does play an important role in our explanations. The Malthusian propensity of populations to grow *rapidly* to environmental limits is one of the processes that generates the competition between variants that in turn drives selection. In essence, at the typical evolutionary time scale we assume that demographic processes acting on a shorter time scale generate some average level of population pressure and that therefore such fast-acting processes are not the rate-limiting steps in the evolutionary process. In a generation-by-generation microevolutionary context such assumptions will be violated and models may have to be adjusted accordingly.

68. Endler's 1986 analysis of the patterns of strength of natural selection in the wild is a nice example of the sorts of generalizations we get in the face of diversity and complexity. Selection is often rather strong, stronger more often than evolutionists' intuitions before Endler's review typically supposed. We also know from the analysis of cross-cultural data (pioneers include Murdock 1949, 1983 and Jorgensen 1980) that cultural variation is not without pattern.

69. The foregoing owes much to the work of Wimsatt 1981.

70. Sober 1991.

71. Ibid., p.

72. Boyd and Richerson 1985; Cavalli-Sforza and Feldman 1981.

73. For the long version of this argument, see Richerson and Boyd 1987.