

Modeling Cultural Evolution

Richard McElreath

Joseph Henrich

January 12, 2006

For publication in *Oxford Handbook of Evolutionary Psychology*, Dunbar and Barret, editors

1 Introduction

When Darwin left for his voyage around the world on the *Beagle*, he took with him the first volume of Charles Lyell's *Principles of Geology*. Later in the voyage he received the second volume by post somewhere in South America. Lyell never accepted Darwin's account of evolution by natural selection, presumably because of his religious beliefs. It is ironic then that Lyell's work played a crucial role in the development of Darwin's thinking. In some ways Lyell's principle of uniformitarianism is as central to Darwinism as is natural selection.

Before Lyell, it was common to explain the features of the earth's geology in terms of past catastrophes: floods, earthquakes and other cataclysms. In contrast, Lyell tried to explain what he observed in terms of the cumulative action of processes that we could observe every day in the world around us—the sinking of lands and the build up of sediments. By appreciating the accumulated small effects of such processes over long time spans, great changes could be explained.

Darwin took the idea of small changes over long time spans and applied it to populations of organisms. Darwin was a good naturalist and knew a lot about the everyday lives of plants and animals. They mate, they reproduce, they move from one place to another,

and they die. Darwin's insight was to see that organisms vary, and the processes of their lives affect which types spread and which diminish. The key to explaining long run change in nature, to explaining the origin of new species, of whole new types of organisms, and of life itself was to apply Lyell's principle of uniformitarianism to populations. By keeping track of how the small events of everyday life change the composition of populations, we can explain great events over long time scales.

Biologists have been thinking this way ever since Darwin, but it is still news in most parts of the social sciences. Are people products of their societies or are societies products of people? The answer must be “both,” but theory in the social sciences has tended to take one side or the other (Marx's dialectic being an obvious exception). In evolutionary models, this classical conflict between explanations at the level of the society (think Durkheimian social facts) and explanations at the level of individuals (think microeconomics) simply disappears. Population models allow explanation and real causation at both levels (and more than two levels) to exist seamlessly and meaningfully in one theory. We don't have to choose between atomistic and group-level explanations. Instead, one can build models about how individuals can create population-level effects which then change individuals in powerful ways.

Cultural evolutionary models are much the same as better-known genetic ones: events in the lives of individuals interact at the scale of populations to produce feedback and powerful long-term effects on behavior. There are three basic steps.

1. One begins by specifying the structure of the population. How large is it? Is it sub-divided? How do sub-divisions affect one another? How does migration work? How is the population size regulated?
2. Then one defines the life cycle of the organism. How does mating work? When is learning possible? What states do individuals pass through from birth to death?
3. Finally, one defines the different heritable variants possible in the model. What is the range of strategies or mutations over which evolution operates? How do these variants affect events in the life cycle of the organism, such as death or development, including learning, attention, and inference?

Since cultural evolutionary models can contain two interacting biological systems of inheritance, culture and genes, the answers to these questions can be different for each system. For example, individuals may be able to acquire many different socially-learned behaviors, but the range of possible genetically inherited learning strategies may be very small. The number of genetic parents has an upper limit of two (for most vertebrates at least), but cultural parents can be many and the contributions among them can be very unequal. In some cultural evolutionary models, the contribution of each parent is typically non-additive in ways most people consider impossible in genetics.

After the structure of the model is completely specified, the objective is to transform

these assumptions into mathematical expressions that tell us how the frequencies of each cultural and genetic variant (and the covariance among them, if necessary) change during each stage of the life cycle. These expressions, called *recursions*, do the work of integrating events in the lives of individuals into micro-evolutionary consequences—changes observable over short time spans. The next goal is to deduce the long-term macro-evolutionary consequences of the assumptions. This is done by finding any combinations of cultural and genetic variants that lead to steady states, *equilibria*, and what combinations of environmental conditions and life-cycle variables make different equilibria possible. Some of these equilibria will be stable, meaning the population will be attracted to them, while others will be unstable, meaning the population will move away from them. Stable equilibria are candidates for long-term evolutionary outcomes, and unstable ones are important because they often inform us as to how likely the population is to reach any of the stable equilibria or how much time it may spend at each.

Thus by writing down formal expressions that capture assumptions about how tiny events in the lives of individuals affect survival, reproduction, and the probabilities of being a cultural parent, evolutionary models allow one to deduce the population-level evolutionary consequences of individual-level psychologies, decision rules, and behavior. At the same time, since these expressions simultaneously define how events in the life-cycle affect the population and how the population affects individuals, it is a two-way street. The mass action arising from individuals integrates up at the population level to have potentially powerful effects on the fates of individuals with different cultural and genetic variants. These different fates in turn lead to further changes in the population, which lead to yet more consequences for individuals.

It is not easy to keep all of these balls in the air simultaneously. The slipperiness of verbal reasoning is famous, and that is perhaps the reason why so many fields, from philosophy to economics to physics, use formalism to make deductions about complex systems. The steady stream of interesting and counter-intuitive results that emerge from these formalisms has demonstrated their value and made them centerpieces of theory development.

Many social scientists and biologists work on how individuals make decisions and how behavior is acquired. Fewer ask how those decisions and mechanisms of learning aggregate at the population level. Our position is that both are inherently interesting and crucial for understanding evolving systems, including culture. In the remainder of this paper, we explore three key, and sometimes controversial, issues in the evolution of culture which arise by examining the population processes cultural inheritance may generate. We invite the reader to join us in a tour of this biological frontier and see how formal population models of cultural systems may clarify and address questions about human behavior, psychology, and society.

2 Why bother with cultural evolution?

Some phenotypes need more than genes and environment, to be represented in a formal model.

Sometimes people ask us why we should even bother with modeling cultural evolution? Why are genetic models not sufficient? What scientific payoff is there in the added complexity?

These are fine questions, and they have fine answers. The basic issue is to identify the minimal requirements for representing evolution of phenotype in a species. For exam-

ple, we could construct a very simple genetic model in which the change (Δ) in the frequency of an allele, p , is a function of environmental state, E . This system would have a single recursion:

$$\Delta p = F(p, E),$$

where the function $F(p, E)$ is to be specified depending upon what model of adaptation to the environment we might choose. It might be that E has little effect on individuals with different alleles, or it might be that E favors one over the others. It might be that E is fluctuating, so that selection favors different alleles at different times. The change might depend upon p itself, as it does in the example of sickle-cell anemia and other cases of overdominance. But nowhere do we allow in such a system for E itself to evolve in response to p .

The scientific question is whether such models are sufficient to model the evolution of a given human phenotype. If we only knew genotypes and the state of the environment, could we predict the behavior of organisms in the next time period? When the answer to this question is “no,” we need at least one more equation:

$$\begin{aligned}\Delta p &= F(p, q, E), \\ \Delta q &= G(p, q, E),\end{aligned}$$

where q is the frequency of some cultural variant (a dialect, say), and $G(p, q, E)$ a function telling us how dialect responds to environment, E , and its own previous state, q , and the frequency of an allele, p .

This all sounds rather complex. And it can be. However, when important parts of phenotype are acquired during development and depend upon previous phenotypes, some system like this is useful for understanding how the organism evolves. Unless we think existing behaviors could be predicted solely from knowing the environment and the distribution of genes, at some point evolutionary

models must incorporate the dynamics of behavioral inheritance. No heroic assumptions are required for behavioral inheritance to exist: if portions of phenotype depend upon the phenotypes of other individuals, then weak or strong inheritance of behavior can exist. In the long run, in a given model, it might turn out that behavioral dynamics have little effect on the outcome. In others, it will make a huge difference.

Cultural evolutionary models (as well as niche construction models, see Odling-Smee et al. 2003) can model just the non-genetic behavioral dynamics, as if q above did not depend upon p , as well as joint dynamics of a coupled gene-culture system (Boyd and Richerson, 1985; Cavalli-Sforza and Feldman, 1981; Durham, 1991). In each case, however, the structure of the model is decided by the question of interest. In the rest of this review, we show how cultural evolution models have been used to address questions about human behavior.

3 Transmission in noisy systems

The imperfection of the analogy between genetic and cultural evolution does not mean culture does not evolve.

While evolutionary principles are equally applicable to almost any dynamical system, many researchers approach models of cultural transmission and evolution via an analogy with genetic evolution. This has led some to be concerned about the strength of this analogy (Sperber, 2000). If cultural variants are not discrete, are prone to “mutation,” and are strongly affected by learning biases, then is it appropriate to speak of “transmission” of culture at all? While we have no particular attachment to the term “transmission,” we think the answer is definitively “yes.” Even

if all the above is true, culture can still be an evolving system that leads to cumulative adaptation. This does not mean that evolved psychology has no role to play in how culture evolves (we think psychology has a huge role to play in understanding culture), but we think it does mean that dismissing cultural evolution on the basis of imperfection of the genetic analogy is unwarranted.

Many people—enthusiasts of the “meme” approach and critics alike—seem to have been persuaded by Richard Dawkins’ abstract statements on what is required for adaptive evolution to occur. In *The Extended Phenotype* (1982), he argued that any successfully replicating entity must exhibit (1) longevity, (2) fecundity, and (3) fidelity. The entity must last long enough (longevity) to make copies of itself (fecundity) that are reasonably similar to it (fidelity). Some have interpreted this to mean that anything with high mutation rates cannot be a successful replicator. Thus if cultural ideas change in the process of social learning, the conclusion is that they do not constitute an evolving system at all (see citations in Henrich and Boyd 2002). Similarly, if cultural variants are continuous and blended entities, then they never exactly replicate, and again cannot produce adaptive evolution.

These conclusions are unfounded. Read very generally, Dawkins’ conditions are necessary and sufficient—There must be some heritability for adaptive evolution to occur. However, there are many ways to produce heritable variation. So in the strict sense many people have read them, while Dawkins’ conditions are sufficient, they are definitely not necessary. Reverse-engineering DNA may tell us how inheritance *can* work, but it does not tell us how it *must* work. Henrich, Boyd and Richerson (forthcoming) examines the problems with this reverse-engineering in greater depth.

In this section, we address concerns arising from the gene-culture analogy by demonstrat-

ing ways that transmission can deviate substantially from the genetic analogy but nevertheless produce both heritable variation and adaptive evolution. Our broader message is that biologists and social scientists alike have tended to think too narrowly in terms of the genes metaphor. Many other systems of inheritance are possible in principle, and culture is only one. We believe that it is more productive to drop the genetic analogy and instead study cultural transmission and evolution on its own.

3.1 Noisy learning can maintain heritable cultural variation

Before the union of genetics and Darwinism, most biologists, including Darwin, thought that inheritance was a blending process: offspring were a mix of parental phenotypes. Darwin was troubled by Fleeming Jenkin’s (1864) argument that natural selection could not produce adaptations, because inheritance would quickly deplete the variation natural selection relies upon. Fisher’s (1918) argument reconciling genetics with continuous phenotypic variation purportedly rescued Darwin, but in reality both Jenkin’s argument and those who think Fisher saved Darwin are simply wrong: blending inheritance can preserve variation, and particulate inheritance is neither necessary nor sufficient to preserve variation (Maynard Smith, 1998, has a chapter that examines this problem).

Boyd and Richerson (1985) presented a simple model to prove this point. The model assumes that (1) cultural variants are continuous (non-discrete), and (2) naive individuals sample n cultural parents and adopt a weighted average of their observed behavior—blending inheritance. Observations and reconstructions are prone to an arbitrary amount of error, however, and therefore inheritance here involves continuous traits, blending, and noise/error. They derive a recursion for the variation in cultural behavior

after one generation of learning. To simplify their presentation, assume that there are only two cultural parents and that each contributes equally to socialization. Let ϵ be the variance in error in cultural learning. When ϵ is large, learning is noisy. When $\epsilon = 0$, cultural variants replicate perfectly. After some calculus, the variation in cultural behavior, V , after learning is (see pages 73-74):

$$V' = \frac{1}{2}(V + \epsilon).$$

If $\epsilon = 0$, then the above has only one stable value, $V' = V = 0$. Blending reduces variation each generation until it is all gone. In this case, Jenkin was correct. However, if $\epsilon > 0$, the equilibrium amount of variation (found where $V' = V$) is:

$$\hat{V} = \epsilon.$$

Thus if there is substantial noise, there will be substantial variation at equilibrium. This variation can be subject to selective forces and produce adaptive change, just as in the genetic case.

The population comes to rest at the amount of variation above, because while blending inheritance does deplete variation, error in learning replenishes it. In the long run, the balance between these two forces results in the expected amount of variation in the population being equal to the average amount of error individuals commit when learning. Blending can never reduce variation below this amount, because as soon as it does, more mistakes are made, and more variation is pumped into the system. Variation cannot stay forever above this amount, because learning averages out any “error” that exceeds that inherent in the learning itself.

Boyd and Richerson also showed that if cultural parents assort by phenotype, then assortment can help to maintain variation. This might occur if similar types inhabit similar environments or if similar types are more likely

to mate and jointly socialize their offspring. When this happens, the parents being blended together are more similar to one another and therefore the loss of variation due to blending is less than in the case above. If parents are weighted unequally (mom is more important than dad), this will also tend to slow the rate at which blending reduces variation, because unequal weighting reduces the effective number of cultural parents.

How cultural learning actually works is a good empirical question, but models like this one prove that the argument that cultural variants cannot evolve in a meaningful way, because they are (1) not discrete entities like genes and (2) prone to error, is simply not a valid deduction. Likewise, the observation that culture does evolve does not imply that there are any units analogous to genes nor that imitation and other forms of social learning are highly accurate.

We also think that the empirical evidence is quite strong that many aspects of human behavior (including technology) evolve in a Darwinian fashion (Richerson and Boyd, 2005). Many of these are not plausibly genetic, in any immediate sense. Thus non-deductive philosophical arguments that culture cannot evolve seem very suspicious, especially when there are existing deductive arguments to the contrary.

3.2 Noisy learning can produce adaptive evolution

Some authors (Sperber and Hirschfeld, 2004) have made a lot out of the results of experiments that resemble games of “telephone” (as it is called in North America) or “Chinese whispers” (as it is called in Britain). When pairs of individuals pass a signal along a chain, the message tends to be corrupted. Thus, we might conclude, social learning is too error prone to maintain variation or content in and of itself. Strong innate inferential mechanisms may be needed to stabilize cultural dif-

ferences, and these cognitive attractors may swamp any evolutionary dynamics possible in culture.

We do not doubt that psychological biases for learning exist, and their importance has long been a part of cultural evolutionary modeling (see Boyd and Richerson’s (1985) direct bias). However, Henrich and Boyd (2002) have addressed whether strong innate inferential mechanisms swamp adaptive cultural evolution—influenced by selective forces like imitating successful people—by deriving a model of cultural transmission that assumes continuously varying representations under the influence of weak selective transmission and strong cognitive attractors. This model addresses the complaint that culturally transmitted ideas are rarely if ever discrete, but instead blend, as well as the complaint that cognitive influences on social learning swamp transmission effects such that cultural variation is not heritable. Using a very general model, they show that these complaints are deductively invalid. In fact, they derive a non-intuitive conclusion: If inferential biases are sufficiently strong relative to selective forces, a continuous representation (quantitative blending) model reduces to the discrete-trait replicator dynamic commonly used in population models of both culture and genes. Thus, powerful and biased inferential mechanisms actually mean that even a weak selective component will eventually determine the final equilibrium of the system, in true Darwinian fashion. The important assumption here is that learners’ psychology have multiple inferential systems—that is, there are multiple cognitive attractors. Strong cognitive biases do not swamp selective effects, but rather make discrete models better estimates of the actual dynamics.

In two other models in the paper, Henrich and Boyd (2002) construct systems with large amounts of transmission error to show that accurate individual-level replication of cultural

variants is not necessary for selective forces to generate either cultural inertia or cumulative cultural adaptation. Their second model shows that if learners aggregate information from multiple cultural parents using a conformist bias (see Henrich & McElreath, this volume) they can dramatically reduce the average noise/error in their inferences, suggesting that our psychology may have genetically evolved a conformist bias, in part, to reduce transmission error (among other reason, see Henrich & Boyd 1998). In the aggregate, this bias creates heritability at population level, and can lead to cultural inertia. In the third model, Henrich and Boyd combine all the potential problems with models of cultural evolution, assuming continuous (non-discrete) cultural representations, incomplete transmission, and substantial inferential transformations. Despite these assumptions, they construct a model which produces adaptive cultural evolution.

3.3 Other inheritance systems

In many baboons, females inherit dominance rank from their mothers and sisters (Silk and Boyd, 1983). In these species, fitness is strongly affected by this extra-genetic inheritance: any female adopted at birth into a high-ranking matriline would be better off than if she were adopted into a low-ranking matriline. And this female will have her dominance rank before she fights a single member of her social group. Dominance is heritable, has important effects on fitness, and yet the mechanism of inheritance is at least partly non-genetic. The rules of how this inheritance works are complicated and very unlike genes. It probably depends upon the composition of ones own matriline, the composition of the entire social group, and local resource density and feeding competition. And yet no primatologist could completely understand baboon biology without taking this complicated extra-genetic pedigree into account. Its exist-

tence may lead females to strive for rank because of its downstream consequences, in addition to its immediate resource access effects (Boyd, 1982; Leimar, 1996).

Extra- or “epigenetic” (Maynard Smith, 1990) systems like this are increasingly recognized: everywhere biologists look, they find hints of inheritance systems either built on top of genes or built from entirely different mechanisms. If the key question is what mechanisms account for heritable phenotypic differences among organisms, then the answer appears to be “many.” Jablonka and Lamb’s *Evolution in Four Dimensions* (2005) mounts the empirically rich argument that heritable differences in many species are due to the action of several inheritance systems (genetic, epigenetic, behavioral and symbolic), sometimes interacting, sometimes acting in parallel.

If one thinks about cell division for a moment, it is obvious that processes other than the replication of DNA are needed to explain how it works. Organelles need to be copied (Sheahan et al., 2004), and the genetic code itself needs to be copied (and this is not contained in the DNA, nor could it be). Beyond cell division, adult phenotypes depend upon imprinting and other forms of learning that may channel the environments offspring are exposed to (a kind of niche construction—Odling-Smee et al. 2003). And finally, most biologists believe that DNA was certainly not the first form of hereditary biological material (Szathmáry and Maynard Smith, 1995). Thus some inheritance systems must be able to sometimes create complementary and even usurping inheritance systems.

In light of these plausible “inheritance systems,” it appears that human culture may not be so special or surprising at all, in the sense of being a non-genetic system of inheritance. Organisms as diverse as arabidopsis (a small plant related to mustard that is a favorite of geneticists), common fruit flies and single-celled microscopic animals such as paramecia

exhibit heritable differences due at least in part to mechanisms other than the sequence of nucleotides in their DNA. The existence of social learning as a system of inheritance and adaptation that functions in complement to DNA may turn out to be unremarkable.

To someone who makes formal models of evolutionary systems, the question that we must answer is whether it will be sufficient to represent human (or any other organism's) evolution with just state variables for its alleles. If we require state variables for early childhood experience, imprinting, or behaviors acquired via social learning, to make useful models of our own evolution, then attempts to construct culture-free models are simply scientifically inadequate. As with each of the possible systems above (e.g. Jablonka and Lamb, 1991; Pál and Miklós, 1999; Maynard Smith, 1990), the specific dynamics and consequences of cultural learning may be rather unique and very important for understanding both micro- and macro-evolution.

In the next two sections, we explore models of the possible dynamic consequences of cultural inheritance. While such models do not tell us how human evolution actually works, they direct our attention to possibilities we are unlikely to consider, if we consider DNA to be the only important source of heritable variation in our species.

4 The relative strength of forces of cultural evolution

Cultural evolution may be most different in the relative difference in strength of evolutionary forces, rather than the absolute speed of its evolution.

It is commonly observed that cultural evolution may be much faster than genetic evolution. Styles of dress and speech, technological

innovations, and reorganizations of human societies happen much faster than the average tempo of genetic evolution. Despite the massive differences in behavior and social organization among human societies, there is little genetic variation among groups within our species (Pääbo, 2001), leading most social scientists to infer that differences among human groups are due to rapid cultural evolution, not selection on genes.

While we agree that cultural evolution is typically absolutely faster than genetic evolution, at least in the short term, this is only part of the story. The danger with the summary we just gave is that it encourages the view that cultural and genetic evolution lead to similar outcomes, only on different time scales. The *relative* rates of competing evolutionary forces are very different in the two systems. Population geneticists tend to think of evolution as the result of the balance of forces acting on alleles. Migration, mutation, and selection all act to alter allele frequencies, but appreciating the balance of these forces is what makes population genetics predictive. Because the balance is likely quite different in cultural models (and presumably the real systems the models caricature), quite different outcomes are possible.

4.1 The balance of selective forces and migration

For our discussion, we focus on the relative strengths of two forces, migration and selection. Selection in the cultural case refers to learning forces that favor some behavioral variants over others. For example, people probably prefer to imitate the successful (see Henrich and McElreath, this volume), and this favors behaviors that lead to success.

An ounce of mixing is a pound of effect, in most models of genetic evolution. In large animals like ourselves, migration among subpopulations is typically a very strong force. This strong force of migration tends to unify

subpopulations of alleles with respect to selection. However this is only true because measured selection coefficients tend to be small, relative to the force of mixing (Endler, 1986). If selection were stronger (and it sometimes is—see again Endler 1986), then more differences could be maintained among sub-groups.

But in a cultural model, the strength of learning biases that, for example, favor behaviors with higher payoffs over behaviors with lower payoffs can be arbitrarily strong. Natural selection of ideas does occur, such as when different fertility ideologies influence the differential growth of religious groups (Stark, 2005). A school of American archaeology used to argue that most important cultural and technological change came about through natural selection of this kind (see Boone and Smith, 1998, for references). It would therefore be useful to consider how strong such selection is, relative to what we might consider fairly fast genetic evolution—such as the 4% increase in the depth of finch beaks Peter and Rosemary Grant recorded on Daphne Major, an island in the Galapagos, during a two-year draught in 1976 and 1977. This strength of selection is sufficient to produce beaks substantially deeper in less than a decade, assuming selection would continue at the same rate (Grant and Grant, 1993).

An extrapolation from an empirical example of cultural evolution will help to make clear how much stronger “selective” forces—by which we mean forces that favor different variants in a non-random way—can be in cultural systems. The classic study of the diffusion of technological innovations is the Ryan and Gross (1943) study of the diffusion of hybrid corn in Iowa farmers. Hybrid corn became available in Iowa in 1928 and was eventually adopted by nearly all farmers by 1941, over a period of 13 years. For those completing and reviewing the study, the shock was how long it took hybrid corn, which had a 20% increase in yield over then-existing va-

rieties, to spread. We want to make something of the opposite point: hybrid corn diffused much more quickly than we might expect, based upon its payoff difference with existing strategies and how natural selection would respond.

If we take the genetic replicator model and use it to model the diffusion of hybrid corn, we can get a feeling for how much stronger selective forces can be in cultural evolution. This thought experiment violates many truths. We are assuming a year is the generation time, and that there is no individual decision-making beyond imitation of successful strategies. However, the ordinary population genetic replicator dynamic and that for simple imitation models is very similar (Gintis, 2000, provides a general derivation). The most basic model, in which individuals compare their own payoff against that of a random individual and preferentially copy the strategy with the higher payoff, yields:

$$\Delta p = p(1 - p)\beta(w_1 - \bar{w}),$$

where p is the frequency of the cultural variant (hybrid corn), β a rate parameter, and w_1 and \bar{w} have similar meanings to the genetic model, payoff to the behavior of interest and the average payoff, respectively.

Figure 1 shows these models with two strengths of “selection,” compared to the actual spread of hybrid corn. At the actual payoff difference between hybrid corn and then-existing varieties, the spread would have been far slower than observed. A difference as large as 50% is needed to predict the diffusion of hybrid corn in 13 years. The actual spread lagged behind this prediction for as much as half of the diffusion period, but then accelerated, so clearly other forces were at work in this example (Henrich, 2001). For current purposes, note that whatever social learning mechanisms are at work here must magnify observed payoff differences. Consider also that a 20% difference in yield is unlikely

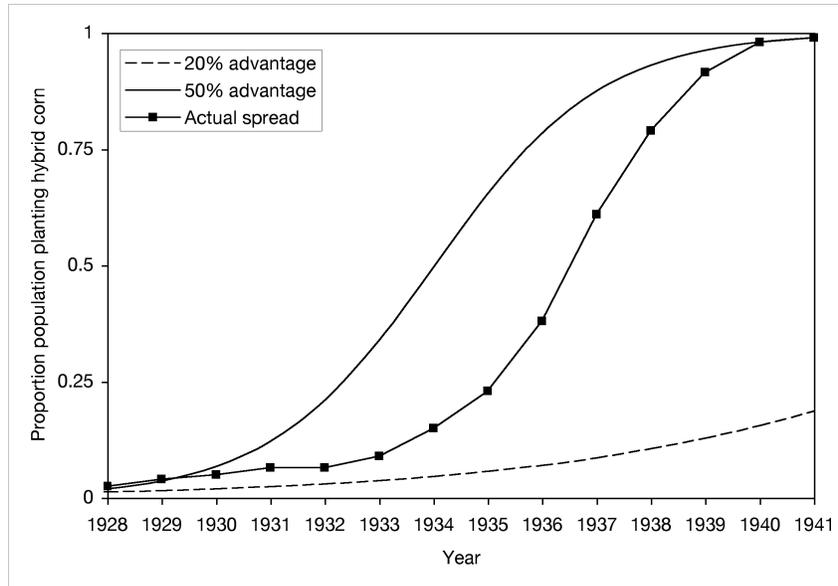


Figure 1: The diffusion of hybrid corn, modeled with the simple replicator dynamic presented in the text, for two strengths of “selection.” The dashed curve shows the predicted spread using the actual payoff difference between hybrid corn and then-existing varieties (20%). If this were natural selection, a 20% difference in fitness would be tremendous and rare, from one generation to the next. At this strength, the curve falls far short of predicting a spread in about 13 years. The solid curve shows the predicted spread for a 50% advantage, which is capable of predicting a spread in the approximately 13 years it took for hybrid corn to diffuse. See Henrich (2001).

to result in a 20% difference in reproduction or survivorship important to natural selection on genes. Many other behaviors matter for aggregate fitness of an individual. Thus the magnitude of “selection” in this case of cultural diffusion seems even larger in comparison to typical genetic estimates.

Because selective forces, arising from human psychology, that favor some variants over others may be strong, and especially strong relative to mixing, cultural evolution may produce outcomes that are very unlikely in genetic evolution. In this section, we explain one important case in which cultural evolutionary models produce equilibrium results that are possible, but highly unlikely, in analogous genetical models. As discussed in Henrich & McElreath (this volume), other examples may include ethnic marking (McElreath et al., 2003) and ethnocentrism (Boyd and Richerson, 1985; Gil-White, 2001).

4.2 Group selection for altruistic behavior

By focusing on competition among cultural groups (cultural group selection), several modeling efforts have demonstrated how the cultural transmission of behaviors related to cooperation and punishment may explain some otherwise puzzling patterns of human prosociality (Boyd and Richerson, 1985; Boyd et al., 2003). See Henrich (2004) and Henrich and Henrich (forthcoming) for reviews. The reason these models can result in stable cooperative equilibria, while analogous genetical models cannot, is due to the plausibility of strong imitation forces opposing forces of mixing (Boyd et al., 2003).

Mixing is an enemy of altruism because selective forces can only produce altruism when the between-group variance in behavior is large enough to overcome within-group selection opposing altruism. Price (1972) and later Hamilton (1975) showed that selection favors

altruism when:

$$\text{var}(p_i)\beta(w_i, p_i) > \text{E}(\text{var}(p_{ij})\beta(w_{ij}, p_{ij})),$$

where p_i is the frequency of an altruism gene in population subdivision i , w_i is the average fitness in group i , and p_{ij} and w_{ij} are the frequency of altruism and fitness of individual j in group i , respectively. $\beta(x, y)$ indicates the slope of the linear regression of x on y ($\partial x / \partial y$). Thus the beta coefficients above are selection gradients for different components of fitness. In plain language, this condition can be read as:

The product of the variance in altruism among groups and the rate of change in the average fitness of individuals in a group as a function of the number of altruists in the group

must exceed

the average product of the variance within each group and the rate of change in individuals fitness as a function of the amount of altruism the individual exhibits.

“Groups” here are defined by the scale at which helping behavior benefits other individuals. Behaviors that aid brothers and behaviors that aid entire ethnic groups are both governed by this fact. However, common descent maintains genetic variation among kin groups, while variation among large groups is much harder to sustain. Mixing is very strong in animals like ourselves, leading to either very little equilibrium variance among large groups of individuals or the steady leaching away of variation (see the model by Rogers, 1990). If learning forces like conformity effectively reduce mixing of cultural variants, then variation among groups can remain high enough to support group selection. There is nothing heretical about this statement. W. D. Hamilton himself saw kin selection as a special case of this general condition (see Hamilton, 1975).

The key issue in any model of the evolution of altruism is what forces are available to maintain variation among groups.

In the cultural case, it is plausible, although hardly yet proven empirically, that strong learning dynamics combined with weak effective migration can result in more variance than analogous genetical models (Boyd et al., 2003, models this process). This in turn might produce selection on culturally transmitted ideas that lead to self-sacrifice. Groups with such ideas may either defeat their neighbors in open conflict, because they can muster more fighters to the field of battle, or defend themselves better from aggression, because they can recruit more people to dig trenches, build walls, or mount a defense of arms.

We must caution the reader to avoid a mistake others have made in understanding this hypothesis. Cultural group selection trades off the very fact that human ethnic groups are well-mixed genetically, but still maintain appreciable cultural distinctiveness. Alleles for group-oriented self-sacrifice are unlikely to spread, because genes move among ethnic and other cultural groups quite often. However, this mixing does not always destroy cultural variation. Immigrants do not necessarily erode the variation in such ideas among groups, because immigrants may quickly conform to local beliefs, even though they cannot change their alleles. The group selection is on culturally transmitted beliefs, not on physical bodies. It is possible to construct a working cultural group selection model (Boyd and Richerson, 2002) in which comparison across groups generates the equilibrium shifts, not differential reproduction or survival of groups. In this case, the group selection may involve no differential death or birth of human bodies at all.

An effect like this might seem initially implausible. Would a system of phenotypic transmission like social learning, created by genetical evolution, actually lead to qual-

itatively different outcomes for an organism? But the evolution of sexual reproduction transformed how traits are inherited and created equally (if not more) novel evolutionary dynamics. Models of sexual selection of animal signals have no problem producing situations in which males produce and females prefer costly ornaments that lower the overall fitness of the population (Fisher, 1930; Lande, 1981). Few people have a problem calling such equilibria fundamentally Darwinian, even though evolution sometimes proceeds quite differently in sexual than asexual species. Similarly, we should not balk at noticing that a genetically-evolved system for acquiring behavior via social learning might end up producing equilibria that are not the self-same ones the genes themselves would be selected to arrive at.

5 Gene-culture coevolution

*Gullibility may be an adaptation,
because critical thinking is costly.*

Over the very long run, cultural dynamics cannot continue to always outrun genetics. Genes must have an eventual influence. One reason could be that, as variation among cultural variants diminishes, the rate of evolution will slow, and then lagging changes in genetic variants will become more important. Also, the cultural system should eventually reach some stationary distribution, even if it is stochastic. Then selection on genes, however slow, may determine how this equilibrium shifts. Even rates of change in classic organic evolution appear to vary on different scales (Penny, 2005). Thus it seems that ignoring genes in the long run is probably a mistake.

In this final section we present a very simple model of gene-culture coevolution. It helps explain one way to model the joint evolution of transmission systems with very different rates of change. Also, this model al-

lows us the opportunity to explain a few important predictions about behavior that arise from gene-culture models.

5.1 When culture is much faster than genes

One way to deal with the difference in rates is to assume that the distribution of cultural variants reaches an equilibrium instantaneously, with respect to genetic evolution. The distribution of alleles then responds to this stationary distribution of cultural variants. Provided cultural dynamics are sufficiently faster than genetic ones, then this method yields a good approximation of the joint system dynamics. Boyd and Richerson (1985) and Alan Rogers (1988) have used this tactic to derive joint evolutionary equilibria for simultaneous cultural and genetic recursions, without resorting to more-complex multi-dimensional techniques. Numerical analysis of the recursions shows that the infinitely-fast-culture assumption does not result in misleading results.

The basic problem is that the change in frequency of a single cultural variant can be represented in a one-dimension system by the abstract function:

$$\Delta p = F(p).$$

This means we can compute the change in the frequency, if we know the current frequency. But if we add a simultaneous second recursion for genes that specify how culture is acquired, then we have a two-dimensional system with two functions:

$$\begin{aligned}\Delta p &= F(p, q), \\ \Delta q &= G(p, q).\end{aligned}$$

Now we must know both the frequency of the cultural variant and the genes influencing social learning in order to find the change in either. The trick is to determine stability in such systems. In principle, stability in

these two-dimensional systems can be solved with matrix techniques. However, if the cultural dynamics are fast enough relative to the genetic dynamics, the cultural dimension p will come to rest at its steady state, \hat{p} , very quickly. This can be true either because there are many opportunities to learn and update behavior per selection event or because selection coefficients are weak, compared to the rate of change due to learning (see the previous section). If either is true, then the system arrives at a cultural equilibrium quickly, and q will respond to this value. As q changes under selection on genes, of course, \hat{p} will also change. But now since p instantly reaches its steady-state for any given value of q , we have a one-dimensional system again:

$$\begin{aligned}\Delta \hat{p} &= F(q), \\ \Delta q &= G(\hat{p}, q).\end{aligned}$$

With such a system, all we have to worry about is the stability of the genetic equilibria. The cultural equilibrium just responds to it. You might think that this means the genes run the show, and that such a model produces the same outcomes as a culture-free model. But as we will demonstrate, not even the simplest models back up that intuition.

Here is a model in the spirit of Rogers (1988). We think this very simple model demonstrates the vulnerability of some commonly held beliefs about what kinds of behavior we expect natural selection to produce. Imagine a simple organism capable of imitating the behavior of older individuals, in addition to investing effort in updating behavior through individual trial and error. We use the discrete formulation, but as with all models of this type, there is an equivalent continuous formulation (in which individuals do some imitation and some individual learning). Each generation, individuals learn according to an inherited allele (individual or social) and then receive payoffs determined by whether what they have learned is adaptive under current

circumstances.

First, a caveat: People sometimes complain that it is unrealistic to consider a pure “social learning” strategy, because real people always make inferences while being influenced by the behavior of others. We agree. All social learning depends upon individual psychology and how that process works. If we expressed this model in its completely equivalent continuous form, with a family of mixed strategies that rely upon a mix of individual and social influence, fewer people would complain. The version we present here is better for illustrating the insights we wish to draw from it. Models are like cartoons: there is an optimal amount of detail, and often that amount is very small. We caution readers of such models not to get hung up on vague words like “social learning” that have different meanings in different sub-disciplines, but instead to attend to the structure of the assumptions. As others have shown, equivalent models can be derived under the assumption that individuals are entirely Bayesian updaters, but able to observe what other people do (Boyd and Richerson, 2005; Bikhchandani et al., 1992).

Suppose an infinite number of behaviors are possible, but only one is adaptive for current environmental circumstances and yields a payoff B . All others yield a payoff of zero. This assumption just sets the scale of payoffs, so we lose no generality with it. The environment itself changes state, making a new behavior optimal, with probability u each generation. When this happens, since there exists a very large number of possible behaviors, we assume all existing behavior in the population is rendered maladaptive. Individual learners pay a cost for experimentation and mistakes (C), but they always arrive at the currently adaptive behavior. In contrast, social learners pay no up-front costs, but they just copy a random adult from the previous generation, so they have no guarantee of acquiring the *currently* adaptive behavior.

With the above assumptions, we can write fitness expressions for each allele, I for individual learners and S for social learners. Let a be the frequency of currently adaptive behavior among adults of the previous generation. Then:

$$\begin{aligned} W(I) &= B - C, \\ W(S) &= Ba. \end{aligned}$$

The variable a is the frequency of adaptive behavior at any one moment, but it changes over time. This implies a recursion for how a changes, and this process will depend upon how the population learns. Let L be the frequency of individual learners in the population. Then the frequency in the next generation is:

$$a' = u(0) + (1 - u)(L(1) + (1 - L)a).$$

The parameter u is the fraction of the time that the environment changes, rendering all past behavior maladaptive, each generation. The rest of the time, L of the population learned for themselves and arrived at adaptive behavior with certainty. The remaining $1 - L$ of the population imitates and transmits the previously adaptive frequency a .

Now we apply the assumption that cultural dynamics are much faster than genetic dynamics. This allows us to find the steady state value of a , call this \bar{a} , for any given L . This exists where $a' = a$ and is:

$$\bar{a} = \frac{L(1 - u)}{1 - (1 - L)(1 - u)}.$$

Over the long run, the fitness of social learners will depend upon this value. We plug \bar{a} into the expression for $W(S)$ and find the value of L that yields a genetic equilibrium, the endpoint of the long-term selection on genes. The equilibrium frequency of individual learning, \hat{L} turns out to be:

$$\hat{L} = \frac{B/C - 1}{1/u - 1}.$$

This expression tells us how the stable frequency of individual learning responds to the

costs of learning and the unpredictability of the environment. The quantity B/C is the ratio of the benefits of acquiring adaptive behavior to the costs of learning it. As this goes down, learning is more costly, and the frequency of individual learning declines. The second effect is that as the environment becomes less predictable (u increases), the denominator decreases and the equilibrium frequency of individual learning increases. If the world is unstable, what your parents did may no longer be adaptive, so it pays more to think for yourself.

The most obvious result of this model is that natural selection can easily favor substantial amounts of social learning. Unless u is very large or B/C is very small, there will be a substantial frequency of social learners at equilibrium.

5.2 Gullibility as an adaptation

An interesting further deduction from the above model is the frequency of adaptive behavior once genes also reach equilibrium. Call this \hat{a} . This is found by substituting the value of \hat{L} for L in the expression for \bar{a} . After simplification:

$$\hat{a} = 1 - \frac{C}{B}.$$

This result is very interesting. Notice that it does not depend upon u . Natural selection adjusts learning in response to u so that, at equilibrium, the value of socially-acquired behavior, \hat{a} , is governed only by the cost of information. When the world is relatively stable from one generation to the next, there are more social learners at equilibrium, which reduces the expected value of socially-acquired behavior. However, the countervailing effect is that, in a more-stable world, adaptive behavior has a better chance to accumulate, so a smaller number of individual learners can provide the same expected accuracy of behavior as a large number, in an unstable world.

Richerson and Boyd (2005) call this effect the *costly information hypothesis*: when information about the world is costly to acquire, it pays to rely upon cheaper ways of learning. Consider what proportion of behavior is adaptive to current circumstances, when C/B is very small, perhaps 1/100. In this case, because information is very cheap to acquire, most individuals (if not all) will be individual learners, and the expected accuracy of behavior, \hat{a} , will be nearly 100%. But when information is costly, because it is dangerous, time-consuming, or difficult to acquire and process, then the expected accuracy will be much smaller.

When we look at a population of animals and ask why they behave as they do, this model (and many others like it, see Boyd and Richerson 1995) suggests it will be risky to assume that development (in this case, learning) is irrelevant to our explanations of what behavior we will see. If the costs of information are high, then substantial portions of the population will be practicing maladaptive behavior.

Moreover, this will be the optimal strategy, from the point of view of the genes. Any more individual learning would not be an equilibrium, even though it would lead to more accurate behavior. What is happening is that social learning saves fitness costs at one point in the life cycle, only to pay other fitness costs later. Even models of cumulative culture (Boyd and Richerson, 1996) show the same tradeoff. When we sample behaviors, we might not notice the information-gathering costs paid by individual learners and conclude that individual learning has higher fitness, because on average individual learners practice more-accurate behavior.

The social learners in this very stylized model are gullible. They believe whatever the previous generation demonstrates. In this case, gullibility is an adaptation, because the costs one would have to pay to verify all the

suggested behavior in the world would be too great. Some individual learning is always favored, because otherwise the population cannot track the environment at all. But large doses of gullibility can be adaptive, because information is costly.

Our impression of real human societies is that many people will believe nearly anything you tell them, at least at first. Many readers of this chapter will be successful students or professional scholars, who have substantial experience with teaching. Isn't it amazing that students are willing to take our word on so many abstruse topics? We think models like this one suggest an answer: being gullible when a problem is abstruse is adaptive, because it is often beyond the individual's means to verify the accuracy of it alone. If we insisted on learning everything for ourselves, we would miss out on many very adaptive solutions. Given how hard it is for agricultural scientists to decide what crops in what proportions to plant, it seems implausible that many real agriculturists, who have to live off their produce, can afford to experiment and analyze their year-to-year yields (Henrich, 2002).

The cost of being adaptively gullible may be that we are sometimes, perhaps often, lead astray. The universal existence of magical thinking might be a symptom of this tradeoff. After all, if you cannot disprove that there are dangerous spirits in the forest, it may be best to just trust that there are. We think this possibility is provocative, because it suggests that behavioral maladaptation in some domains may be a standard feature of human societies, even in our proper evolutionary niche. Clearly some of the non-adaptive things humans do, like gorge on fatty foods, are probably products of our minds being adapted to a world in which fatty foods were rare. But this kind of out-of-equilibrium explanation of maladaptive behavior is not the only possibility. If the above model captures the texture of human cultural adaptations, then even in Pleis-

tocene environments, human societies probably showed widespread non-adaptive and maladaptive behaviors, because such behaviors are a side-effect of relying upon social learning as an adaptation. Thus we suspect that even Pleistocene foragers did a lot of "silly" things, from the perspective of genetic fitness. This would not mean that culture, as a system of behavioral evolution, is a maladaptation, no more than the cost of producing sons means sexual reproduction is a maladaptation. But it would unfortunately mean that predicting behavior is harder than simply applying optimality criteria. It would mean that both development and population dynamics are central to Darwinian explanations of individual behavior.

References

- Bikhchandani, S., Hirshleifer, D., and Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100:992–1026.
- Boone, J. and Smith, E. (1998). Is it evolution yet? A critique of evolutionary archaeology. *Current Anthropology*, 39:141–173.
- Boyd, R. (1982). Density-dependent mortality and the evolution of social interactions. *Animal Behaviour*, 30:972–982.
- Boyd, R., Gintis, H., Bowles, S., and Richerson, P. J. (2003). The evolution of altruistic punishment. *Proceedings of the National Academy of Sciences USA*, 100:3531–3535.
- Boyd, R. and Richerson, P. (1996). Why culture is common, but cultural evolution is rare. *Proceedings of the British Academy*, 88:77–93.
- Boyd, R. and Richerson, P. J. (1985). *Culture and the Evolutionary Process*. Univ Chicago Press, Chicago.

- Boyd, R. and Richerson, P. J. (1995). Why does culture increase human adaptability? *Ethology and Sociobiology*, 16:125–143.
- Boyd, R. and Richerson, P. J. (2002). Group beneficial norms spread rapidly in a structured population. *Journal of Theoretical Biology*, 215:287–296.
- Boyd, R. and Richerson, P. J. (2005). Rationality, imitation, and tradition. In *The Origin and Evolution of Cultures*, chapter 18, pages 379–396. Oxford University Press, Oxford.
- Cavalli-Sforza, L. L. and Feldman, M. W. (1981). *Cultural transmission and evolution: a quantitative approach*. Princeton University Press, Princeton.
- Dawkins, R. (1982). *The Extended Phenotype: The Long Reach of the Gene*. Oxford University Press, Oxford.
- Durham, W. H. (1991). *Coevolution: Genes, Culture, and Human Diversity*. Stanford University Press, Stanford.
- Endler, J. (1986). *Natural Selection in the Wild*. Princeton University Press, Princeton.
- Fisher, R. A. (1918). The correlation between relatives on the supposition of Mendelian inheritance. *Transactions of the Royal Society of Edinburgh*, 52:399–433.
- Fisher, R. A. (1930). *The genetical theory of natural selection*. Dover, New York.
- Gil-White, F. J. (2001). Are ethnic groups biological 'species' to the human brain?: Essentialism in our cognition of some social categories. *Current Anthropology*, 42:515–554.
- Gintis, H. (2000). *Game Theory Evolving*. Princeton University Press, Princeton.
- Grant, B. R. and Grant, P. R. (1993). Evolution of Darwin's finches caused by a rare climatic event. *Proceedings of the Royal Society London B*, 251:111–117.
- Hamilton, W. D. (1975). Innate social aptitudes of man: an approach from evolutionary genetics. In Fox, R., editor, *Biosocial Anthropology*, pages 133–153. Malaby Press, London.
- Henrich, J. (2001). Cultural transmission and the diffusion of innovations: Adoption dynamics indicate that biased cultural transmission is the predominate force in behavioral change and much of sociocultural evolution. *American Anthropologist*, 103:992–1013.
- Henrich, J. (2002). Decision-making, cultural transmission and adaptation in economic anthropology. In Ensminger, J., editor, *Theory in Economic Anthropology*, pages 251–295. AltaMira Press.
- Henrich, J. (2004). Cultural group selection, coevolutionary processes and large-scale cooperation. *Journal of Economic Behavior and Organization*, 53:3–35.
- Henrich, J. and Boyd, R. (2002). On modeling cognition and culture: Why replicators are not necessary for cultural evolution. *Journal of Cognition and Culture*, 2:87–112.
- Henrich, J., Boyd, R., and Richerson, P. J. (forthcoming). Five common mistakes in cultural evolution. In Sperber, D., editor, *Epidemiology of Ideas*. Open Court Publishing.
- Henrich, N. and Henrich, J. (forthcoming). The Origins of Cooperation: A cultural and evolutionary exploration among the Iraqi Chaldeans of metro-Detroit. Oxford UP, New York.

- Jablonka, E. and Lamb, M. (2005). *Evolution in Four Dimensions*. MIT Press, Massachusetts.
- Jablonka, E. and Lamb, R. M. (1991). Sex chromosomes and speciation. *Proceedings of the Royal Society of London, B*, 243:203–208.
- Jenkin, F. (1864). The origin of species. *North British Review*, 46:277–318.
- Lande, R. (1981). Models of speciation by sexual selection on polygenic traits. *Proceedings of the National Academy of Sciences USA*, 78:3721–3725.
- Leimar, O. (1996). Life history analysis of the Trivers-Willard sex-ratio problem. *Behavioral Ecology*, 7:316–325.
- Maynard Smith, J. (1990). Models of a dual inheritance system. *Journal of Theoretical Biology*, 143:41–53.
- Maynard Smith, J. (1998). *Evolutionary Genetics*. Oxford University Press, Oxford.
- McElreath, R., Boyd, R., and Richerson, P. J. (2003). Shared norms and the evolution of ethnic markers. *Current Anthropology*, 44:122–129.
- Odling-Smee, F. J., Laland, K. N., and Feldman, M. W. (2003). *Niche Construction: The Neglected Process in Evolution*. Number 37 in Monographs in Population Biology. Princeton University Press, Princeton.
- Pääbo, S. (2001). The human genome and our view of ourselves. *Science*, 16:1219–1220.
- Pál, C. and Miklós, I. (1999). Epigenetic inheritance, genetic assimilation and speciation. *Journal of Theoretical Biology*, 200:19–37.
- Penny, D. (2005). Relativity for molecular clocks. *Nature*, 436:183–184.
- Price, G. R. (1972). Extension of covariance selection mathematics. *Annals of Human Genetics*, 35:485–490.
- Richerson, P. J. and Boyd, R. (2005). *Not by Genes Alone: How culture transformed human evolution*. University of Chicago Press, Chicago.
- Rogers, A. R. (1988). Does biology constrain culture? *American Anthropologist*, 90:819–831.
- Rogers, A. R. (1990). Group selection by selective emigration: The effects of migration and kin structure. *American Naturalist*, 135:398–413.
- Ryan, B. and Gross, N. (1943). The diffusion of hybrid seed corn in two iowa communities. *Rural Sociology*, 8:15–24.
- Sheahan, M. B., Rose, R. J., and McCurdy, D. W. (2004). Organelle inheritance in plant cell division: the actin cytoskeleton is required for unbiased inheritance of chloroplasts, mitochondria and endoplasmic reticulum in dividing protoplasts. *The Plant Journal*, 37:379–390.
- Silk, J. B. and Boyd, R. (1983). Female cooperation, competition, and mate choice in matrilineal macaque groups. In Wasser, S. K., editor, *Social Behavior of Female Vertebrates*, pages 315–347. Academic Press, New York.
- Sperber, D. (2000). An objection to the memetic approach to culture. In Aunger, R., editor, *Darwinizing Culture: The Status of Memetics as a Science*, pages 163–173. Oxford University Press, Oxford.
- Sperber, D. and Hirschfeld, L. A. (2004). The cognitive foundations of cultural stability and diversity. *Trends in Cognitive Sciences*, 8:40–46.
- Stark, R. (2005). *The rise of Mormonism*. Columbia University Press, New York.
- Szathmáry, E. and Maynard Smith, J. (1995). The major evolutionary transitions. *Nature*, 374:227–232.