

Original Article

A Bayesian approach to the evolution of social learning

Charles Perreault^{a,*}, Cristina Moya^b, Robert Boyd^{a,b}

^a*Santa Fe Institute, Santa Fe, NM 87501, USA*

^b*Department of Anthropology, University of California, Los Angeles, Los Angeles, CA 90095, USA*

Initial receipt 3 June 2011; final revision received 20 December 2011

Abstract

There has been much interest in understanding the evolution of social learning. Investigators have tried to understand when natural selection will favor individuals who imitate others, how imitators should deal with the fact that available models may exhibit different behaviors, and how social and individual learning should interact. In all of this work, social learning and individual learning have been treated as alternative, conceptually distinct processes. Here we present a Bayesian model in which both individual and social learning arise from a single inferential process. Individuals use Bayesian inference to combine social and nonsocial cues about the current state of the environment. This model indicates that natural selection favors individuals who place heavy weight on social cues when the environment changes slowly or when its state cannot be well predicted using nonsocial cues. It also indicates that a conformist bias should be a universal aspect of social learning.

© 2012 Elsevier Inc. All rights reserved.

Keywords: Cultural evolution; Learning psychology; Conformism; Bayesian learning; Bet-hedging; Origin of culture

1. Introduction

Social learning is an essential part of human adaptation and is likely a key factor generating our remarkable ecological success over the last 50,000 years (Richerson & Boyd, 2005). Social learning has been observed in a wide range of other species in diverse taxa including mammals (Galef & Laland, 2005; Perry & Manson, 2003), birds (Benskin et al., 2002; Lefebvre, 2000), fish (Brown & Laland, 2003), and even invertebrates (Leadbeater & Chittka, 2007).

There has been much interest in understanding the evolution of social learning (Boyd & Richerson, 1985; Cavalli-Sforza & Feldman, 1981; Kameda & Nakanishi, 2003; McElreath & Strimling, 2008; Rendell, Fogarty et al., 2010; Rogers, 1988; Whiten et al., 2001). Investigators have tried to understand when natural selection will favor individuals who imitate others, rather than learning on their own. They have also tried to understand how selection shapes the process of imitation. What happens if there are a number of potential models exhibiting different behavior?

How should observable characteristics of models, such as indicators of fitness, affect the imitation process?

In all of this work, social learning and individual learning are treated as alternative, conceptually distinct processes. Social learning is conceived as a transmission process in which the determinants of behavior are transmitted socially from one individual to another. This transmission process may be subject to errors, biases, and systematic transformations, but most work assumes that social learning leads to reasonably accurate copying. Then, to build models of cultural evolution, investigators modify mathematical models drawn from population genetics or epidemiology to account for the novel features of social learning. Most important among these are that (1) behaviors that are acquired or modified by individual learning can subsequently be transmitted and (2) social learning can be biased so that some variants are more likely to be transmitted than others. These processes are modeled as deviations from accurate, unbiased transmission. Often, it is assumed that natural selection determines the relative importance of social and individual learning so as to maximize genetic fitness. This work has been widely influential, transforming the idea of cultural evolution from a vague analogy to an active area of both theoretical and empirical research (Mesoudi, 2011).

* Corresponding author. Tel.: +1 505 946 2768; fax: +1 505 982 0565.
E-mail address: cperreault@santafe.edu (C. Perreault).

A number of authors have criticized this approach to the evolution of social learning on the grounds that social learning and individual learning are not psychologically distinct processes (Heyes, 1994; Plotkin, 1988). Indeed, both individual learning and social learning involve cue-based inferences about what is the best behavior in the organism's environment. Other authors have complained that much of the theoretical literature has assumed that social learning and individual learning are alternatives competing for determination of phenotype when in fact they are usually complementary processes that lead individuals in the same direction (Laland, 2004).

Here we present a simple model in which both individual learning and social learning are modeled as arising from a single learning process. We assume that learning can be modeled as Bayesian inference. This provides a useful framework for studying learning and cognitive development (Perfors & Tenenbaum et al., 2011). In this case, the environment varies, and the adaptive problem is to infer the current state of the environment using two sources of information: the behavior of the individuals from the previous generation (social cues) and information about the current state of the environment that is learned individually, such as through a trial-and-error process (environmental or nonsocial cues). There is no assumption that behavior is transmitted or copied. Rather, the central adaptive problem faced by individuals is how to behave given the observed social and environmental cues. Answering this question is tricky because the usefulness of the social cue depends on how individuals in previous generations combined social and environmental cues. When environmental cue allows accurate inference about the current environment, social learning might not be needed. Returning to home base empty-handed for several days might be a good indication that game is rare in the region and that hunting there is not the best idea. However, many adaptive problems are difficult because the environment does not provide clear cues to the best behavior. What is the best design for a bow? What causes malaria? It is not clear what decision rule will be favored by selection when the environmental cue does not allow accurate inference. The problem is that the quality of the social cue depends on the long-term effects of how individuals in the population integrate social and environmental information in their decisions. Thus, to determine the optimal reliance on social cues, it is necessary to model the coevolution of the culturally transmitted pool of information and the genes that determine how this information is transmitted. This problem is further complicated by the fact that these genes will respond to selection on individuals, not to the effect of the average quality of information on the population as a whole.

Below, we derive the evolutionary stable learning rule that specifies how much weight individuals should put on social information given some environmental cue. We find that (1) a reliance on imitation is favored when individual

learning is inaccurate and environments are not too variable; (2) social learning increases average fitness because it allows individual learning to be restricted to situations in which it is accurate; and (3) when learners can observe the behavior of three or more individuals from the previous generation, they should show a conformist bias, that is, they should place a disproportionate weight on the more common behavior.

2. The model

2.1. A learning model with two cues

Consider a large population that lives in an environment that has two states: *state 1* and *state 2*. Each generation, the environment switches from the state that it is in to the alternate state with probability γ and stays in the same state with probability $1-\gamma$. Thus, over the long run, the environment is equally likely to be in each state. Individuals acquire one of two behaviors: *behavior 1* and *behavior 2*. Individuals exhibiting behavior 1 have fitness $1+d$ when the environment is in state 1 and have fitness 1 when the environment is in state 2. Similarly, behavior 2 has fitness $1+d$ when the environment is in state 2 and has fitness equal to 1 when it is in state 1. Thus, individuals need to determine the current state of the environment in order to choose the favored behavior. Individuals with the favored behavior have higher reproductive success, transmitting their genotype at a higher rate to the next generation.

Individuals have access to two cues that provide information about the current state of the environment. We assume that both cues can be represented by numbers and that the values of the cues observed by a given individual are x and y . Let $\Pr(x|1)$ and $\Pr(x|2)$ be the probability that an individual observes cue value x in environments 1 and 2, respectively. Similarly, let $\Pr(y|1)$ and $\Pr(y|2)$ be the probability that an individual observes cue value y in environments 1 and 2. Using Bayes law, the conditional probability that the environment is in state 1 given the cue values x and y is (see Electronic Supplementary Material for derivation, available on the journal's website at www.ehbonline.org):

$$\Pr(1|x, y) = \frac{\Pr(x|1)\Pr(y|1)}{\Pr(x|1)\Pr(y|1) + \Pr(x|2)\Pr(y|2)} \quad (1)$$

The conditional probability that the environment is in state 2 is just $1-\Pr(y|1)$.

Due to the symmetry of the model, an organism maximizes expected fitness by choosing the behavior that is best in the environment that is most likely given the observed cues. If state 1 is more likely to be the current state $\Pr(1|x, y) > 1/2$, choose behavior 1; if state 2 is more likely, $\Pr(1|x, y) < 1/2$, choose behavior 2. Thus, to maximize expected fitness, individuals should choose behavior 1 when the joint probability of the observed cues given

environmental state 1 is higher than the joint probability given environmental state 2, or

$$Pr(x|1)Pr(y|1) > Pr(x|2)Pr(y|2) \quad (2)$$

Note that, in a temporally varying environment, selection need not favor behavioral rules that maximize expected fitness. Instead, selection will favor behavioral rules that maximize geometric mean fitness, and when selection is strong, this can lead to “bet hedging” rules that sacrifice expected fitness in order to reduce the variance in fitness (Gillespie, 1977).

Learning requires information about the predictive values of the two cues. In a Bayesian framework, this information is represented as the conditional probabilities of different cue values given different environments. While Bayesian models are good approximations for some kinds of learning (Kalish et al., 2007; Tenenbaum et al., 2006), it is unlikely that organisms actually carry out the necessary calculations. Rather, this framework gives a simple representation of how prior information should be combined with cue observations to generate expected fitness maximizing inferences about the environment. The information about the predictive value of the different cues could come from a previous episode of learning or be innate knowledge created by natural selection.

2.2. A model of social learning

So far, we have not made any assumptions about the nature of the cues—they are simply observations that individuals make and that can be used to predict the current environmental state. To model social learning, we assume that individuals have access to one nonsocial and one social cue. The nonsocial cue, x , is something that the organism observes in the environment. Values of this “environmental” cue are normally distributed with mean μ and variance $v=1$ when the environment is in state 1 and with mean $-\mu$ and variance $v=1$ when the environment is in state 2. Thus, positive values of x support the inference that the environment is in state 1, and negative values support the inference that it is in state 2. The parameter μ is a measure of how accurately the environmental cue allows individuals to predict the state of the environment. The predictive value of the environmental cues increases with μ .

The social cue is obtained when the organism observes the behavior of n individuals sampled at random from the population in the previous generation. Following the social learning literature, we refer to these individuals as “models.” The value of the social cue, y , is the number of observed models exhibiting behavior 1. Let p be the expected frequency of the behavior favored by selection in the population. Thus, p is the frequency of behavior 1 when the environment is in state 1, and from the symmetry of the model, $1-p$ is the frequency of behavior 1 when the environment is in state 2. Then $Pr(y|1)$ is binomial with

parameters p and n , and $Pr(y|2)$ is binomial with parameters $1-p$ and n .

With these assumptions, Eq. (2) says to adopt behavior 1 if

$$y - \frac{1}{2}n > -gx \quad (3)$$

where

$$g = \frac{\mu/v}{\ln\left(\frac{p}{1-p}\right)} \quad (4)$$

(See Electronic Supplementary Material for the derivation, available on the journal’s website at www.ehbonline.org.) When the expected frequency of the favored behavior, p , is greater than one half, g is positive. Thus, if x is positive and $y \geq \frac{1}{2}n$, both environmental and social cues indicate that the environment is in state 1, and the condition is always satisfied. Similarly, if x is negative and $y \leq \frac{1}{2}n$, the condition is never satisfied. When the cues point in opposite directions, their magnitudes become important. Suppose, for example, that an individual observes more than half of her models using behavior 1 but observes an environmental cue that indicates she is in environment 2 (a negative value of x). Then her choice of behavior depends on the relative magnitude of the cues and the parameter g . If $y - \frac{1}{2}n > -gx$, the optimal decision is to adopt behavior 1. If this condition is not satisfied, she should adopt behavior 2. Thus, for a given pair of cue values, this decision depends on the parameter g that summarizes the relative predictive value of the two kinds of cues. If g is small, the environmental cue is not as good a predictor as the social cue so that it is weighted less heavily; if it is large, the environmental cue is a better predictor. If g were the correct value, an individual following the rule given in Eq. (3) would maximize expected fitness. However, this parameter depends on the expected frequency p and the mean and variance of the cue distribution, μ and v , which are features of the environment that the individuals cannot observe directly.

We assume that g is a heritable attribute of the organism’s psychology that is shaped by natural selection. Individuals can only observe particular cue values y and x , but not p or μ . However, natural selection can shape the learning mechanisms so that they incorporate this information. To model this process, we assume that the value of g is affected by one of m alleles at a haploid locus. The i th allele has a learning rule characterized by the parameter g_i and has frequency S_i . Individuals first acquire their genotype through genetic transmission. Then they observe members of the previous generation and an environmental cue, and determine whether they should adopt behavior 1 or behavior 2. Finally, viability selection adjusts the genotypic frequencies.

The fitnesses of the alleles depend on the average frequency of the two behaviors in the population over time. When the adaptive behavior is common, selection will favor small values of g that place a heavy emphasis on the social cue; when it is not, selection will favor larger values of g . This means that we must also derive an expression for how these frequencies change in the population as a result of learning, given the distribution of different values of g in the population. The distributions of learning rules and behaviors coevolve, each affecting the dynamics of the other.

To model this process, suppose that the environment is in state 1, so that the favored behavior is 1. We denote that the frequency of behavior 1 observed in the previous generation is q . When the environment is in state 1, $q=p$; otherwise, $q=1-p$.

Then the probability that an individual samples y models with behavior 1 from the previous generation, $B(y|q,n)$, is binomial with parameters q and n . The probability that an individual that carries allele i and samples y models exhibiting behavior 1 adopts behavior 1 is $1 - F_1\left[\left(\frac{n}{2} - y\right)\frac{1}{g_i}\right]$, where F_1 is the cumulative normal distribution with mean μ and variance 1. Thus, the probability that an individual with genotype i acquires behavior 1 in environment 1, $\psi_1(g_i)$, is

$$\psi_1(g_i) = \sum_{y=0}^n \left(1 - F_1\left[\left(\frac{n}{2} - y\right)\frac{1}{g_i}\right]\right) B(y|q,n)$$

The frequency of genotype i is S_i , and thus the frequency of individuals with behavior 1 after learning, q' is

$$q' = \sum_i S_i \psi_1(g_i)$$

Similarly, when the environment is in state 2

$$q' = \sum_i S_i \psi_2(g_i)$$

where

$$\psi_2(g_i) = \sum_{y=0}^n \left(1 - F_2\left[\left(\frac{n}{2} - y\right)\frac{1}{g_i}\right]\right) B(y|q,n)$$

and F_2 is the cumulative normal distribution with the mean $-\mu$ and variance 1.

Once every individual has adopted a behavior, natural selection occurs. If the environment is in state 1, the frequency of allele i in the next generation is

$$S'_i = S_i \left(\frac{1 + d\psi_1(g_i)}{\bar{W}}\right)$$

where d is a fitness advantage associated with the favored behavior and $\bar{W} = \sum_i (1 + d\psi_1(g_i))$ is the average fitness of all alleles. Similarly, when the environment is in state 2

$$S'_i = S_i \left(\frac{1 + d\psi_2(g_i)}{\bar{W}}\right)$$

Finally, the state of the environment changes between each generation with probability γ and stays the same with a probability $1-\gamma$.

3. Results

We have explored how natural selection affects g by numerically iterating this system of recursions. We have studied a total of 4000 different alleles under selection. The values of the g alleles under selection ranged from $g_1=0.005$ to $g_{4000}=20$, with increments of 0.005. At the start of each simulation, the frequency of each allele was $1/4000$; the environment was in state 1; and p , the frequency of the favored behavior in the population, was set to 1. The fitness benefits associated with the favored behavior, d , was set to 0.5. We used two criteria to determine the steady-state genotypic frequencies: either one of the alleles has reached a frequency of 0.999 or the allele with modal frequency remains the same for 1,000,000 generations. Once this steady state is reached, the frequencies of the two behaviors within the population continue to fluctuate, but these fluctuations do not substantially affect genotypic frequencies because rates of genetic change are slow enough that they average over the environmental fluctuations.

We varied the following parameters: the number of models whose behavior can be observed, n ; the absolute value of the mean of environmental cue values, μ ; and the rate at which the environment changes, γ . The number of models observed, n , was set to 1, 3, 8, and 16. The absolute value of the mean of the environmental cue distribution, μ , was varied from 0.1, 0.5 and 1. Note that as μ gets smaller, the accuracy of the environmental information decreases. When $\mu=0$, the environment provides no information about its state since the chance of adopting the favored behavior via individual learning is 0.5—the equivalent of flipping a coin to make a decision. Hence, the probability of getting the better behavior through individual learning alone is 0.54, 0.69, and 0.84 when $\mu=0.1, 0.5,$ and 1, respectively. The probability that the environment changes between each generation, γ , was 0.001, 0.01, and 0.1.

3.1. Slowly changing environments and inaccurate individual learning favor more social learning

Fig. 1 shows the probability of acquiring the favored behavior (assuming that the environment is in state 1) as a function of the frequency of that behavior in the population (horizontal axes) at the evolutionary steady

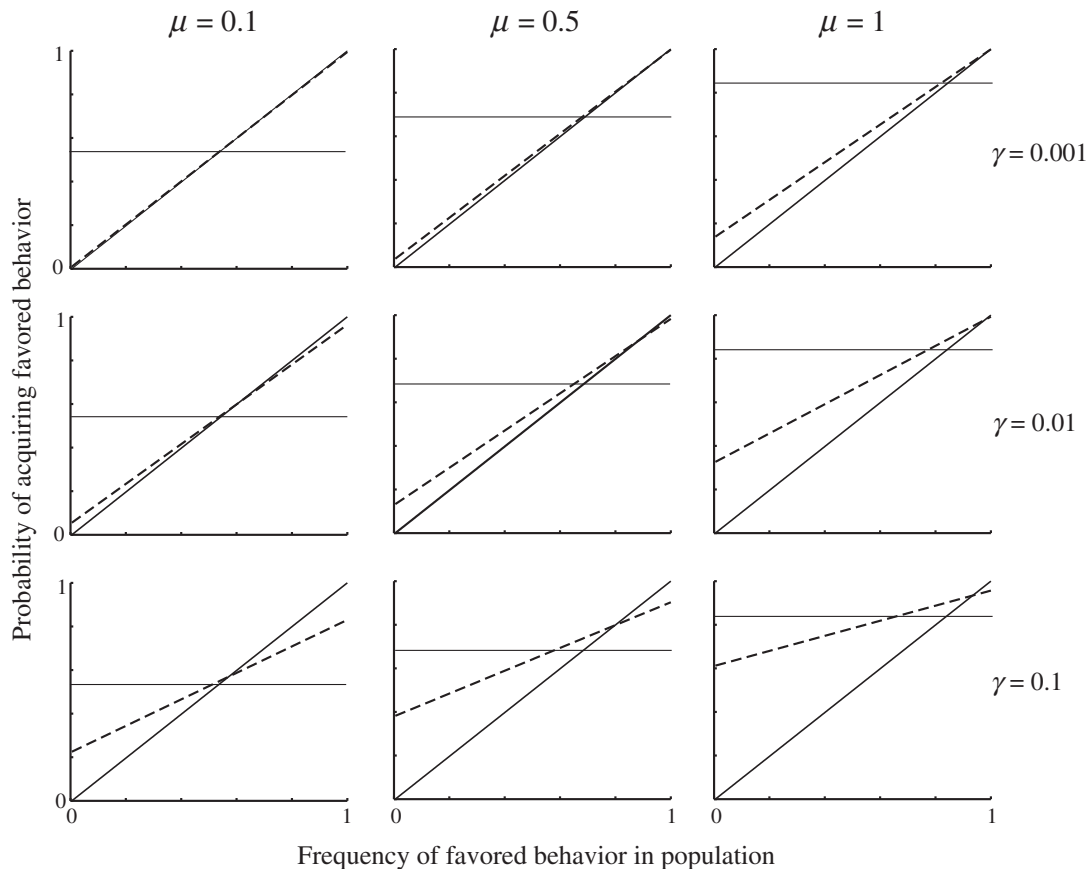


Fig. 1. Probability of acquiring the favored behavior as a function of the frequency of the behavior in the population, p , at the equilibrium steady state (dotted line) when individuals observe $n=1$ model for three levels of environmental cue quality ($\mu=0.1, 0.5, 1$) and three different rates of environmental change ($\gamma=0.001, 0.01, 0.1$), assuming that learners are only able to observe the behavior of one model ($n=1$). These plots show how the frequency of the favored behavior in the population affects the evolutionary dynamics of the cultural traits given that g is at its equilibrium value. We use these plots to summarize the evolutionary equilibria rather than the equilibrium values of g because the plots are easier to interpret. It is useful to compare the evolutionary equilibrium outcome to two benchmarks (also plotted). First, the horizontal line gives the probability of acquiring the favored behavior if the learner completely ignores the social information and instead relies only on the environmental cue—pure individual learning. Second, the 45° identity line running from the origin to the upper right-hand corner gives the probability of acquiring the favored behavior if the learner ignores the environmental cue and instead imitates a random individual from the previous generation—pure unbiased social learning. In this model, the equilibrium

state for three levels of environmental cue quality ($\mu=0.1, 0.5, 1$) and three different rates of environmental change ($\gamma=0.001, 0.01, 0.1$), assuming that learners are only able to observe the behavior of one model ($n=1$). These plots show how the frequency of the favored behavior in the population affects the evolutionary dynamics of the cultural traits given that g is at its equilibrium value. We use these plots to summarize the evolutionary equilibria rather than the equilibrium values of g because the plots are easier to interpret.

It is useful to compare the evolutionary equilibrium outcome to two benchmarks (also plotted). First, the horizontal line gives the probability of acquiring the favored behavior if the learner completely ignores the social information and instead relies only on the environmental cue—pure individual learning. Second, the 45° identity line running from the origin to the upper right-hand corner gives the probability of acquiring the favored behavior if the learner ignores the environmental cue and instead imitates a random individual from the previous generation—pure unbiased social learning. In this model, the equilibrium

learning rule always mixes individual and social learning. However, we can judge the importance of the two cues in determining behavior by asking how close the evolutionary equilibrium rule is to these two benchmarks.

The results plotted in Fig. 1 indicate that inaccurate environmental information and slowly changing environments favor a strong reliance on social cues. Unsurprisingly, when the environmental information is accurate, natural selection favors individuals who put a lot of weight on it. As a result, for any value of γ , the equilibrium learning rule recursion lies closer to individual learning when μ is large (Fig. 1). Conversely, as the quality of the environmental information decreases, selection favors putting less weight on it and relying instead on social information. Similarly, for any value of μ , the equilibrium learning rule recursion lies closer to the individual learning recursion as γ gets larger and the environment becomes less stable. We were thus able to obtain results that are consistent with previous models of the evolution of social learning (Boyd & Richerson, 1985, 1987).

3.2. When social learning is important and more than one model is available, selection favors a conformist bias

Fig. 2 plots the evolutionary equilibrium learning rule when more than one model is available to imitate, for the same range of parameter values as Fig. 1. It shows that as n increases, individuals are increasingly likely to adopt the most common cultural behavior, as long as $p > 1/2$. This phenomenon has been labeled conformist-biased social learning. For example, suppose that a naïve individual observes 10 models; six exhibit one behavior and four the alternative. A conformist individual will adopt the most common behavior with a probability greater than 0.6 (Boyd & Richerson, 1982, 1985; Henrich & Boyd, 1998). This contrasts with unbiased imitators who each copy one randomly selected individual and, as a result, will adopt the most common behavior with probability 0.6 (Boyd & Richerson, 1982, 1985). The effect of conformism can be seen for all parameter combinations. However, the positive effect of conformist bias increases when individual learning is more accurate (μ is large) and when rates of environmental change are low (low γ). These are the conditions under

which the population will spend more of its time with high frequencies of the favored trait such that the benefits of conformism can be captured while avoiding its pitfalls when the adaptive cultural trait is rare. Finally, it is also interesting that the incremental effect of increasing n on conformism decreases as n increases. For example, the difference between the equilibrium learning rule when $n=1$ and $n=3$ is greater than the difference between $n=3$ and $n=8$, and this is greater than the difference between $n=8$ and $n=16$.

A conformist bias creates an obvious problem. When the favored behavior is common, paying attention to the common type is adaptive, but when the favored behavior is rare, conformity is maladaptive. This problem is especially important immediately after the environment changes when most models will have the wrong behavior. Fig. 3 illustrates how natural selection, by acting on g , can optimally make this trade-off. The recursion plot 3b shows the population dynamics when all individuals have the evolutionarily stable learning psychology (i.e., the equilibrium value of g). Notice that the probability of acquiring the favored behavior always exceeds the frequency of the favored behavior in the population. This means that the

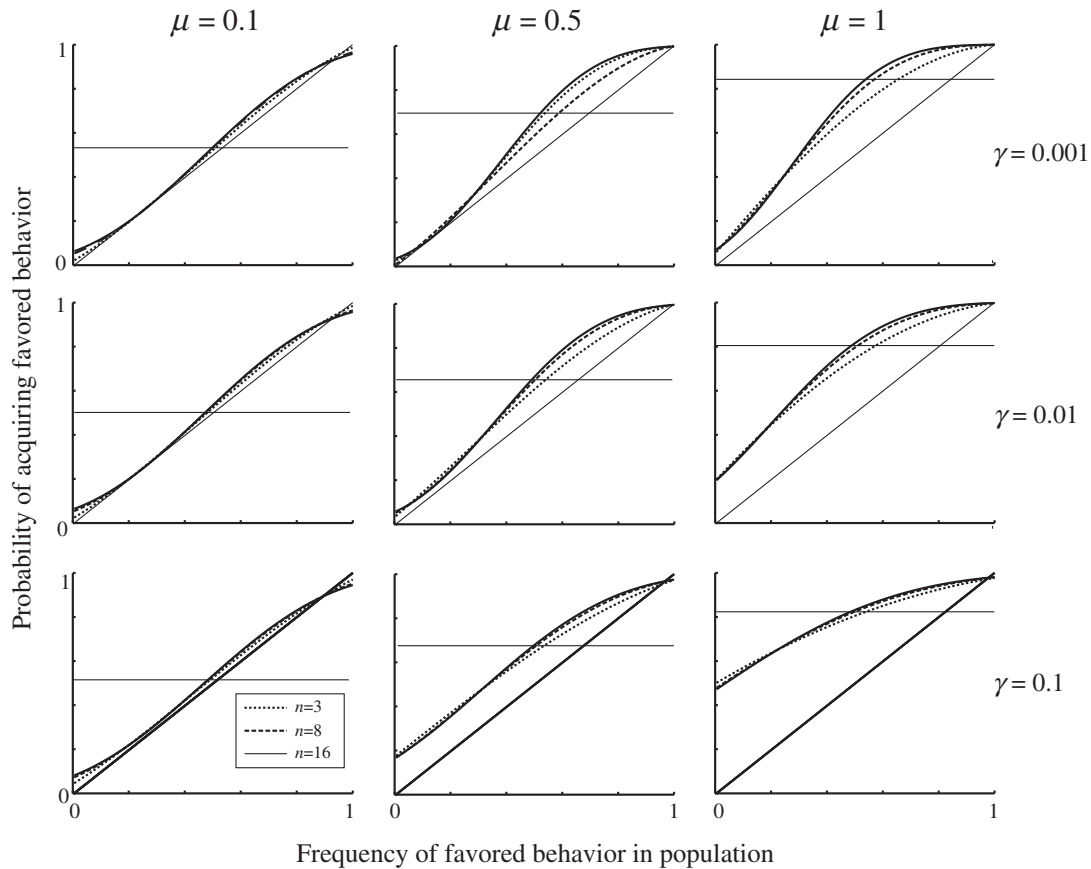


Fig. 2. Probability of acquiring the favored behavior as a function of the frequency of the behavior in the population, p , at the equilibrium steady state when individuals observe $n=3$, $n=8$, or $n=16$ models for three levels of environmental cue quality (μ) and three different rates of environmental change (γ) assuming that the error rate in assessing model's behavior (ϵ) is zero. The 45° identity line gives the probability of acquiring the favored behavior when learners ignore the environmental cue and imitate a random individual from the previous generation (pure unbiased social learning). Note that the probability of acquiring the favored behavior is equivalent to the frequency of the favored behavior in the population during the next generation (p').

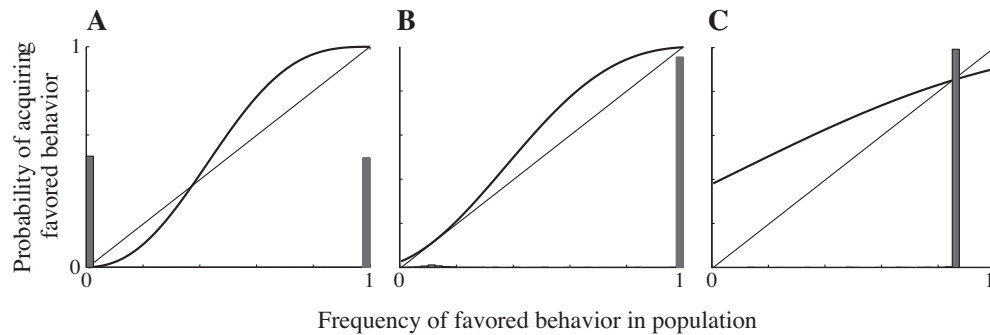


Fig. 3. Probability of acquiring the favored behavior as a function of the frequency of the behavior, p , for values of g that are lower (A, $g=1$) or higher (C, $g=5$) than the equilibrium value (B, $g=1.66$). The superimposed histograms show the proportion of time steps spent by the population at the various frequencies of the favored behavior. Note that all individuals in the population have the same value of g . Recursions are for the following parameter set: $n=8$, $\mu=0.5$, and $\gamma=0.001$.

frequency of the favored behavior always increases through time (unless the environment shifts). The rate of increase is very slow when it is rare; nonetheless, most of the time, the favored behavior has very high frequency ($p=.95-1.0$), but there is a small amount of time its frequency is around 0.1. Lower g values result in individuals that weigh social information more heavily, and as a result, the favored behavior cannot increase after an environmental change, as is the case in Fig. 3A. When this happens, the population spends half the time with the favored behavior as the most frequent one and half the time with the favored behavior almost absent. On the other hand, g values higher than the equilibrium value cause an individual to weigh environmental information more heavily. As a result, the favored behavior increases more rapidly after an environmental change, but it also results in a lower stable equilibrium of p (Fig. 3C). The equilibrium learning rule is the one that allows the population to maximize the time spent at high p values (i.e., high frequencies of the favored behavior).

Our results suggest that because of this trade-off, selection is maximizing expected fitness when $n=1$, but bet hedging when $n>1$. If selection shapes g so as to maximize expected fitness, the equilibrium steady state (ESS value of g should equal the value given by Eq. (4) calculated using the observed steady-state frequency of the favored behavior p . If the ESS value of g is larger than the predicted value, individuals are not maximizing expected fitness, but instead are relying more on the environmental cue. Since the environmental cue is noisy, this leads to bet hedging—increasing the chance of picking the wrong behavior so as to increase the chance of the lineage having some individuals carrying the best variant after an environmental switch. The patterns produced for ESS values of g are consistent with this interpretation. Substituting the actual steady-state values of the frequency of the favored behavior into Eq. (4) yields a predicted value of g assuming that selection is maximizing expected fitness. (Fig. 4) plots the difference between the observed ESS values of g/n and the value of g/n predicted assuming expected fitness maximization. When $n=1$, the predicted and computed values are the same on average, suggesting that selection is maximizing

expected fitness. However, when $n>1$, the ESS values of g/n are systematically larger than the predicted values. In general, selection maximizes geometric mean fitness in a temporally varying environment, but in our simulations, the difference in g values between competing alleles is very small, so selection is weak, and therefore, geometric mean fitness is approximately equal to expected fitness. When $n>1$, however, even small differences in g can lead to big differences in geometric mean fitness. As the number of models observed increases, conformism potentially becomes more costly as it can trap a population at the maladaptive behavior (low values of p). Such deleterious effect of conformism on geometric mean fitness does not exist when $n=1$, as there can be no conformism with one model. As a result, we believe that selection adjusts g so as to create the highest long-term frequency of the favored behavior consistent with not being permanently caught with the wrong behavior.




Fig. 4. The difference between the ESS value of g and the expected value of g based on expected fitness maximization divided by n . The data points represent different combinations of the parameters μ and γ . When $n=1$, the average over a number of runs is zero, indicating that selection is maximizing expected fitness. When $n>1$, the average difference is positive, indicating that selection favors a heavier reliance on the environmental cue than would be predicted by expected fitness maximization.

3.3. Social learning increases average fitness

In the simplest models of the evolution of cultural transmission, social learning evolves but does not increase average fitness, a phenomenon sometimes labeled “Rogers Paradox” (Rogers, 1988). Average fitness increases if models and learners are related (Lehmann et al., 2010; Rendell, Boyd et al., 2010; Rendell, Fogarty et al., 2010) or if social learning allows for the selective use of environmental information (i.e., use environmental information when it is accurate; otherwise, imitate) (Boyd & Richerson, 1995; Enquist et al., 2007).

Selective learning is the essence of the Bayesian model we analyze here and is consistent with the latter argument that social learning increases average fitness in this model. Fig. 5 shows the average fitness advantage of the various evolutionary equilibria, relative to the expected fitness of pure individual learners, as a function of μ and γ . The relative advantage of social learning is larger when the

quality of the environmental cue is poor (when μ is low). The fitness advantage associated with social learning is greater in a stable environment (low γ), as higher γ increases the likelihood that the social information in a population is outdated and inaccurate.

Interestingly, increasing the number of models sometimes decreases average fitness. Fig. 5 shows how the number of individuals observed (n) impacts the relative average fitness advantage of social learning. When the environment is unstable ($\gamma=0.1$, Fig. 5A), expected fitness increases with n . In more stable environments ($\gamma=0.01$), the fitness advantage of social learning as a function of n is concave with a peak at $n=3$ (Fig. 5B). This suggests that when the environment is stable and environmental cue is noisy, there are an optimal number of models, and paying attention to more models decreases expected fitness. We believe that this is an artifact that results from the fact that, in the current model, learners have a very simple prior belief about the long run average frequency of the favored

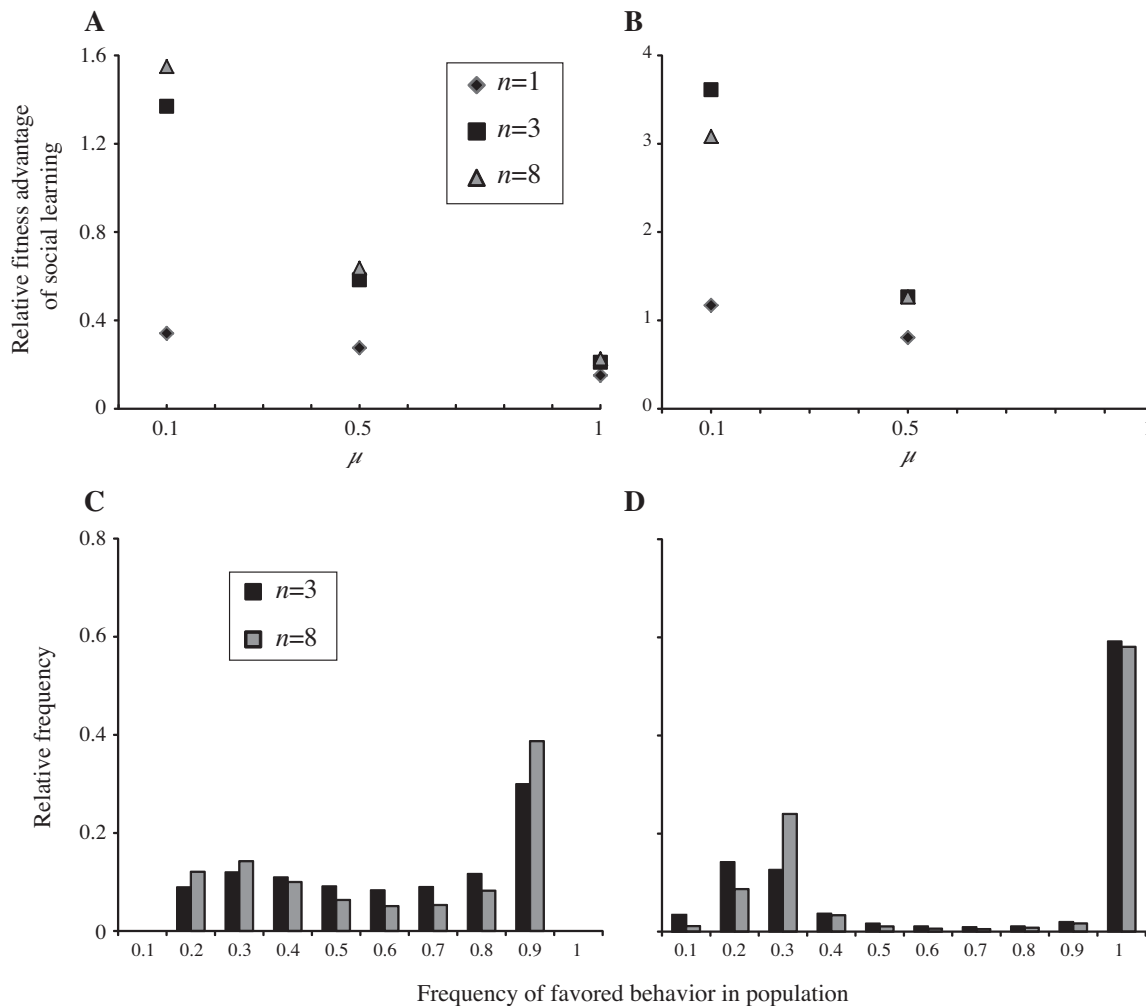


Fig. 5. Fitness of the evolutionary equilibrium learning rules relative to the expected fitness of pure individual learners for different values of n when (A) $\gamma=0.1$ and (B) $\gamma=0.01$. The histograms (C) and (D) show the relative frequency of time spent by the population at different values of the frequency of the favored behavior p for values of γ corresponding to (A) and (B).

variant. As a result, the learning psychology has only one parameter that adjusts both the extent to which individuals rely on social cues and, given that they do, the degree of conformism. Reducing the number of parents effectively reduces the amount of conformism but has a smaller effect on the reliance on social cues, and as a result, the population spends less time at low frequencies of the favored behavior and almost the same amount of time at high frequencies as shown in Fig. 5D. We conjecture that this effect would disappear if actors were endowed with a more flexible prior belief system.

3.4. Imitation error or migration increases reliance on social information when more than one model is available

In previous work, it has been argued that both spatial variation in fitness (McElreath et al. (in press)) and inaccurate imitation (Henrich & Boyd, 2002) favor an increased the reliance on conformism. Here we describe a modification of the model to allow for spatial variation or imitation error. Assume with probability $1-\varepsilon$ individuals accurately observe the behavior of a model, but with probability ε , the learner believes that the model has behavior 1 50% of the time, and behavior 2 50% of the time. This can be thought of as due to a failure to accurately perceive the model's behavior or due to an influx of migrants from a large population in which $p=.5$. This would be the case in an infinite island model in which the environments in different islands fluctuated independently. Under these conditions, a migration rate of 0.1 would mean that, 10% of the time, a learner would learn from a migrant, who is expected to have behavior 1 50% of the time. Thus, the population-level consequence of a migration rate of 0.1 is equivalent to a rate of observation failure of 0.1.

Adding imitation error or migration ($\varepsilon=0.1$) decreases the accuracy of social information—the observed behavior of others is a poorer predictor of the environmental state (see Table 1 in the Electronic Supplementary Material, available online on the journal's website at www.ehbonline.org). When learners can only observe one model ($n=1$), this leads to an increased reliance on environmental information. Surprisingly, however, when individuals have access to more than one model ($n>1$), adding imitation error or migration actually *decreases* the equilibrium value of g and thus increases the reliance on social cues. This interesting result reflects the fact that increasing g also increases the degree of conformism, and this allows learners to circumvent the deleterious effect of imitation error when $p>.5$. By observing more than one individual, learners can conform to the majority of models, which is in general still accurately perceived since the magnitude of imitation errors is not sufficiently large to make the disfavored behavior the dominant one when the favored behavior is dominant and vice versa. However, average population fitness decreases when imitation errors are added.

4. Discussion

The model of the evolution of social learning analyzed above is based on two assumptions. The first one is that natural selection shapes learning mechanisms so that alternative cues about the environment are weighted optimally according to Bayes' law so that reliance on a cue depends on the extent to which that cue predicts fitness-relevant contingencies. The second assumption is that the predictive value of other individuals' observed behavior depends on the past history of social learning and environmental change. This means that determining the evolutionarily stable social learning psychology requires modeling both the coevolution of the pool of culturally transmitted information and the genes that determine how individual psychology processes this information.

Taking this approach, we have shown that selection favors a psychology that causes individuals to rely heavily on imitating the behavior of others when environmental cues are not very informative and environments change slowly. We also show that cultural transmission errors and migration from other environments increase reliance on imitation when learners can observe the behavior of more than one cultural model. These results replicate previous work on the topic (Boyd & Richerson, 1985, 1987; Henrich & Boyd, 1998; McElreath et al. in press) without making the assumption that social learning and individual learning are distinct processes. This suggests that the qualitative conclusions drawn from these studies and this paper are robust and do not hinge on the specificity of any particular learning mechanism.

We have also found that a conformist bias is likely to be a universal aspect of social learning consistent with some previous work (Boyd & Richerson, 1982, 1985; Henrich & Boyd, 1998). Understanding the conditions under which conformist-biased social learning evolves is important because it could be a human psychological adaptation, and that can maintain stable between-group cultural differences which create the conditions for group selection, allowing large-scale cooperation and altruistic punishment to evolve (Boyd & Richerson, 1985; Henrich & Boyd, 1998; Richerson & Boyd, 2005).

Recently, there has been debate about the conditions under which conformism will evolve. Eriksson et al. (2007), Wakano and Aoki (2007), McElreath et al. (in press), Andrés Guzmán and Rodríguez-Sickert 2007. Henrich and Boyd (1998) analyzed models similar to ours in which the environment changes spatially and temporally between two states. Through individual learning or cultural transmission, individuals adopt one of two behaviors, each of which is favored in one of the environmental states. They show that conformism is favored over unbiased imitation by natural selection under a broad range of environmental conditions. Erikson, et al. (2007) have criticized these models on the ground that they assume that individuals have inherent knowledge of the entire set of possible cultural variants

existing in their population and that this set is limited to only two variants. They modify Henrich and Boyd's (1998) model to relax these assumptions, and argue that conformists often do worse than unbiased learners. Wakano and Aoki (2007) have reanalyzed Henrich and Boyd's (1998) model and concluded that conformism can evolve but under a narrower range of environmental conditions than the first study argues. They explain that this discrepancy results from the fact that Henrich and Boyd (1998) did not let their simulations run to convergence. However, Wakano and Aoki (2007) did not include spatial variation in their model, and as a result, it is unclear to what extent the two studies can be compared. Finally, McElreath et al. (in press), stressing the importance of testing the social strategies against both temporally and spatially varying environments, have found that conformism-biased learning yields higher fitness than unbiased social learning as long as the migration rate between populations is not too high. This result is even more pronounced as the cost of individual learning increases. On the other hand, they found that temporal variation favors less conformism, but that spatial variation mitigates this effect.

The current model may explain why different investigators have derived different results about the evolution of conformism. One of the main limitations of the previous models is that they do not account for the fact that the evolution of conformism is likely to depend on the number of models observed. Both Henrich and Boyd (1998) and McElreath et al. (in press) fix this number at 3, while Eriksson et al. (2007) set it to 10. It is easy to see that a majority among three models (e.g., two out of three individuals display behavior R) is not as likely to be informative about the state of the environment as a majority among 10 individuals (e.g., 7 out of 10). The Bayesian approach we used in this paper allows individuals to adjust the extent to which they are conformist depending on the number of models they observe. Overall, our results largely support the idea that conformist individuals are favored by natural selection under a wide range of environmental conditions. We have found that some level of conformism is favored by natural selection even when rates of temporal variation are high. What is more, increasing the number of models observed leads to a greater reliance on conformism and thus broadens the range of environmental conditions that will favor conformism. It is reasonable to assume that during the course of evolution of our species, individuals routinely have had access to many models, which raises the possibility that the current studies are underestimating the range of environmental conditions that favor conformism.

Bayesian methods are particularly well suited to the study of the evolution of cultural evolution. We believe that cultural transmission, as an adaptation, has been designed to solve an inference problem: given some environmental and social information, how should an individual infer the state of its current local environment in order to adopt the behavior that will maximize its fitness? Learning about the

state of the world, whether through an individual trial-and-error process or through social imitation, can be conceived of as a Bayesian process. This study shows that Bayesian tools can be profitably applied to the study of the evolution of social learning.

Supplementary Materials

Supplementary data to this article can be found online at doi:10.1016/j.evolhumbehav.2011.12.007.

Acknowledgments

The authors thank the California Department of Motor Vehicles for queues of eureka-moment-inducing length. We also thank Richard McElreath, Tom Griffiths, and Magnus Enquist for useful comments on earlier drafts of this manuscript. C.P. was supported by the Social Sciences and Humanities Research Council of Canada (752-2006-2301), C.M. and R.B. are supported by NIH grant number 1RC1TW008631.

References

- Andrés Guzmán, R., & Rodríguez-Sickert, C., et al. (2007). When in Rome, do as the Romans do: the coevolution of altruistic punishment, conformist learning, and cooperation. *Evolution and Human Behavior*, 28, 112–117.
- Benskin, C. M. H., & Mann, N. I., et al. (2002). Social learning directs feeding preferences in the zebra finch, *Taeniopygia guttata*. *Animal Behavior*, 64, 823–828.
- Boyd, R., & Richerson, P. (1982). Cultural transmission and the evolution of cooperative behavior. *Human Ecology*, 10, 325–351.
- Boyd, R., & Richerson, P. J. (1985). *Culture and the evolutionary process*. London: University of Chicago Press.
- Boyd, R., & Richerson, P. J. (1987). The evolution of ethnic markers. *Cultural Anthropology*, 2, 65–79.
- Boyd, R., & Richerson, P. J. (1995). Why does culture increase human adaptability? *Ethology and Sociobiology*, 16, 125–143.
- Brown, C., & Laland, K. N. (2003). Social learning in fishes: a review. *Fish and Fisheries*, 4, 280–288.
- Cavalli-Sforza, L. L., & Feldman, M. V. (1981). *Cultural transmission and evolution: a quantitative approach*. Princeton: Princeton University Press.
- Enquist, M., & Eriksson, K., et al. (2007). Critical social learning: a solution to Rogers's Paradox of nonadaptive culture. *American Anthropologist*, 109, 727–734.
- Eriksson, K., & Enquist, M., et al. (2007). Critical points in current theory of conformist social learning. *Journal of Evolutionary Psychology*, 5, 67–87.
- Galef, B. G., & Laland, K. N. (2005). Social learning in animals: empirical studies and theoretical models. *Bioscience*, 55, 489–499.
- Gillespie, J. H. (1977). Natural selection for variances in offspring numbers: a new evolutionary principle. *The American Naturalist*, 111, 1010–1014.
- Henrich, J., & Boyd, R. (1998). The evolution of conformist transmission and the emergence of between-group differences. *Evolution and Human Behavior*, 19, 215–241.
- Henrich, J., & Boyd, R. (2002). On modeling cognition and culture: why replicators are not necessary for cultural evolution. *Journal of cognition and culture*, 2, 87–112.
- Heyes, C. M. (1994). Social learning in animals: categories and mechanisms. *Biological Review*, 69, 207–232.
- Kalish, M. L., & Griffiths, T. L., et al. (2007). Iterated learning: intergenerational knowledge transmission reveals inductive biases. *Psychonomic Bulletin & Review*, 14, 288–294.

- Kameda, T., & Nakanishi, D. (2003). Does social/cultural learning increase human adaptability? Rogers's question revisited. *Evolution and Human Behavior*, 24, 242–260.
- Laland, K. N. (2004). Social learning strategies. *Learning & Behavior*, 32, 4–14.
- Leadbeater, E., & Chittka, L. (2007). Social learning in insects—from miniature brains to consensus building. *Current Biology*, 17, R703–R713.
- Lefebvre, L. (2000). Feeding innovations and their cultural transmission in bird populations. In C. Heyes, & L. Huber (Eds.), *The evolution of cognition* (pp. 311–328). London: MIT Press.
- Lehmann, L., & Feldman, M. W., et al. (2010). Cumulative cultural dynamics and the coevolution of cultural innovation and transmission: an ESS model for panmictic and structured populations. *Journal of Evolutionary Biology*, 23, 2356–2369.
- McElreath, R., Fasolo, B., et al. (in press). The evolutionary rationality of social learning In R. Hertwig and U. Hoffrage, MIT Press.
- McElreath, R., & Strimling, P. (2008). When natural selection favors imitation of parents. *Current Anthropology*, 49, 307–316.
- Mesoudi, A. (2011). *Cultural evolution*. Chicago: University of Chicago Press.
- Perfors, A., & Tenenbaum, J. B. (2011). A tutorial introduction to Bayesian models of cognitive development. *Cognition*, 220, 301–321.
- Perry, S. E., & Manson, J. H. (2003). Traditions in monkeys. *Evolutionary Anthropology: Issues, News, and Reviews*, 12, 71–81.
- Plotkin, H. C. (1988). Learning and evolution. In H. C. Plotkin, Ed. *The role of behavior in evolution* (pp. 133–164). Cambridge, MA: MIT Press.
- Rendell, L., & Boyd, R., et al. (2010). Why copy others? Insights from the social learning strategies tournament. *Science*, 328, 208–213.
- Rendell, L., & Fogarty, L., et al. (2010). Rogers' paradox recast and resolved: population structure and the evolution of social learning strategies. *Evolution*, 64, 534–548.
- Richerson, P. J., & Boyd, R. (2005). *Not by genes alone: how culture transformed human evolution*. Chicago: University of Chicago Press; 2005.
- Rogers, A. R. (1988). Does biology constrain culture? *American Anthropologist*, 90, 819–831.
- Tenenbaum, J. B., & Griffiths, T. L., et al. (2006). Theory-based Bayesian models of inductive learning and reasoning. *Trends in Cognitive Sciences*, 10, 309–318.
- Wakano, J. Y., & Aoki, K. (2007). Do social learning and conformist bias coevolve? Henrich and Boyd revisited. *Theoretical Population Biology*, 72, 504–512.
- Whiten, A., & Hinde, R. A., et al. (2001). Discussion meeting issue 'culture evolves'. *Philosophical Transactions of the Royal Society B-Biological Sciences*, 366.