



ELSEVIER

Evolution and Human Behavior xx (2012) xxx–xxx

**Evolution
and Human
Behavior**

Original Article

A Bayesian approach to the evolution of social learning

Charles Perreault^{a,*}, Cristina Moya^b, Robert Boyd^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 learning and social learning arise from a single learning process. Individuals use Bayesian inference 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.

© 2011 Published by Elsevier Inc.

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.). 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 171

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

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

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

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

197 2.2. A model of social learning

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

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

parameters p and n , and $Pr(y|2)$ is binomial with parameters
225 $1-p$ and n . 226

227 With these assumptions, Eq. (2) says to adopt behavior
228 1 if 228

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

229 where 229

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

230 (See Electronic Supplementary Material for the derivation, 230
231 available on the journal’s website at www.ehbonline.org.) 235

236 When the expected frequency of the favored trait, p , is
237 greater than one half, g is positive. Thus, if x is positive
238 and $y \geq \frac{1}{2}n$, both environmental and social cues indicate
239 that the environment is in state 1, and the condition is
240 always satisfied. Similarly, if x is negative and $y \leq \frac{1}{2}n$, the
241 condition is never satisfied. When the cues point in
242 opposite directions, their magnitudes become important.
243 Suppose, for example, that an individual observes more
244 than half of her models using behavior 1 but observes an
245 environmental cue that indicates she is in environment 2 (a
246 negative value of x). Then her choice of behavior depends
247 on the relative magnitude of the cues and the parameter g .
248 If $y - \frac{1}{2}n > -gx$, the optimal decision is to adopt
249 behavior 1. If this condition is not satisfied, she should
250 adopt behavior 2. Thus, for a given pair of cue values, this
251 decision depends on the parameter g that summarizes the
252 relative predictive value of the two kinds of cues. If g is
253 small, the environmental cue is not as good a predictor as
254 the social cue so that it is weighted less heavily; if it is
255 large, the environmental cue is a better predictor. If g were
256 the correct value, an individual following the rule given in
257 Eq. (3) would maximize expected fitness. However, this
258 parameter depends on the expected frequency p and the
259 mean and variance of the cue distribution, μ and v , which
260 are features of the environment that the individuals cannot
261 observe directly. 261

262 We assume that g is a heritable attribute of the
263 organism’s psychology that is shaped by natural selection.
264 Individuals can only observe particular cue values y and x ,
265 but not p or μ . However, natural selection can shape the
266 learning mechanisms so that they incorporate this informa-
267 tion. To model this process, we assume that the value of g
268 is affected by one of m alleles at a haploid locus. The i th
269 allele has a learning rule characterized by the parameter g_i
270 and has frequency S_i . Individuals first acquire their
271 genotype through genetic transmission. Then they observe
272 members of the previous generation and an environmental
273 cue, and determine whether they should adopt behavior 1
274 or behavior 2. Finally, viability selection adjusts the
275 genotypic frequencies. 275

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

287 To model this process, suppose that the environment is in
 288 state 1, so that the favored behavior is 1. We denote that the
 289 frequency of behavior 1 observed in the previous generation
 290 is q . When the environment is in state 1, $q=p$; otherwise,
 291 $q=1-p$.

292 Then the probability that an individual samples y
 293 models with behavior 1 from the previous generation,
 294 $B(y|q, n)$, is binomial with parameters q and n . The
 295 probability that an individual that carries allele i and
 296 samples y models exhibiting behavior 1 adopts behavior 1
 297 is $1 - F_1\left[\left(\frac{g}{2} - y\right)\frac{1}{g_i}\right]$, where F_1 is the cumulative normal
 298 distribution with mean μ and variance 1. Thus, the
 299 probability that an individual with genotype i acquires
 300 behavior 1 in environment 1, $\psi_1(g_i)$, is

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

303 The frequency of genotype i is S_i , and thus the frequency of
 Q4304 individuals with behavior 1 after learning, q' is

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

306 Similarly, when the environment is in state 2

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

308 where

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

313 and F_2 is the cumulative normal distribution with the mean
 314 $-\mu$ and variance 1.

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

$$319 S'_i = S_i \left(\frac{1 + d\psi_1(g_i)}{-W}\right)$$

where d is a fitness advantage associated with the favored 320
 behavior and $-W = \sum_i (1 + d\psi_1(g_i))$ is the average 321
 fitness of all alleles. Similarly, when the environment is 322
 in state 2 323

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

Finally, the state of the environment changes between each 324
 generation with probability γ and stays the same with a 327
 probability $1-\gamma$. 328

3. Results 329

We have explored how natural selection affects g by 330
 numerically iterating this system of recursions. We have 331
 studied a total of 4000 different alleles under selection. The 332
 values of the g alleles under selection ranged from $g_1=0.005$ 333
 to $g_{4000}=20$, with increments of 0.005. At the start of each 334
 simulation, the frequency of each allele was $1/4000$; the 335
 environment was in state 1; and p , the frequency of the 336
 favored behavior in the population, was set to 1. The fitness 337
 benefits associated with the favored trait, d , was set to 0.5. 338
 We used two criteria to determine the steady-state genotypic 339
 frequencies: either one of the alleles has reached a frequency 340
 of 0.999 or the allele with modal frequency remains the same 341
 for 1,000,000 generations. Once this steady state is reached, 342
 the frequencies of the two behaviors within the population 343
 continue to fluctuate, but these fluctuations do not 344
 substantially affect genotypic frequencies because rates of 345
 genetic change are slow enough that they average over the 346
 environmental fluctuations. 347

We varied the following parameters: the number of 348
 models whose behavior can be observed, n ; the absolute 349
 value of the mean of environmental cue values, μ ; and the 350
 rate at which the environment changes, γ . The number of 351
 models observed, n , was set to 1, 3, 8, and 16. The absolute 352
 value of the mean of the environmental cue distribution, μ , 353
 was varied from 0.1, 0.5 and 1. Note that as μ gets smaller, 354
 the accuracy of the environmental information decreases. 355
 When $\mu=0$, the environment provides no information about 356
 its state since the chance of adopting the favored behavior via 357
 individual learning is 0.5—the equivalent of flipping a coin 358
 to make a decision. Hence, the probability of getting the 359
 better behavior through individual learning alone is 0.54, 360
 0.69, and 0.84 when $\mu=0.1, 0.5$, and 1, respectively. The 361
 probability that the environment changes between each 362
 generation, γ , was 0.001, 0.01, and 0.1. 363

3.1. Slowly changing environments and inaccurate 364 individual learning favor more social learning 365

Fig. 1 shows the probability of acquiring the favored 366
 trait (assuming that the environment is in state 1) as a 367
 function of the frequency of that trait in the population 368
 (horizontal axes) at the evolutionary steady state for three 369




Fig. 1. Probability of acquiring the favored trait as a function of the frequency of the trait in the population, p , at the equilibrium steady state (dotted line) when individuals observe $n=1$ model for three levels of environmental cue quality (μ) and three different rates of environmental change (γ). The horizontal line gives the probability of acquiring the favored trait when learners ignore social information (i.e., pure individual learning). The 45° identity line gives the probability of acquiring the favored trait 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 trait is equivalent to the frequency of the favored trait in the population during the next generation (p').

370 levels of environmental cue quality ($\mu=0.1, 0.5, 1$) and
 371 three different rates of environmental change ($\gamma=0.001,$
 372 $0.01, 0.1$), assuming that learners are only able to observe
 373 the behavior of one model ($n=1$). These plots show how the
 374 frequency of the favored trait in the population affects the
 375 evolutionary dynamics of the cultural traits given that g is
 376 at its equilibrium value. We use these plots to summarize
 377 the evolutionary equilibria rather than the equilibrium
 378 values of g because they are easier to interpret.

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

learning rule always mixes individual and social learning. 390
 However, we can judge the importance of the two cues in 391
 determining behavior by asking how close the evolutionary 392
 equilibrium rule is to these two benchmarks. 393

The results plotted in Fig. 1 indicate that inaccurate 394
 environmental information and slowly changing environ- 395
 ments favor a strong reliance on social cues. Unsurpris- 396
 ingly, when the environmental information is accurate, 397
 natural selection favors individuals who put a lot of weight 398
 on it. As a result, for any value of γ , the equilibrium 399
 learning rule recursion lies closer to individual learning 400
 when μ is large (Fig. 1). Conversely, as the quality of the 401
 environmental information decreases, selection favors 402
 putting less weight on it and relying instead on social 403
 information. Similarly, for any value of μ , the equilibrium 404
 learning rule recursion lies closer to the individual learning 405
 recursion as γ gets larger and the environment becomes 406
 less stable. We were thus able to obtain results that are 407
 consistent with previous models of the evolution of social 408
 learning (Boyd & Richerson, 1985, 1987). 409

410 3.2. When social learning is important and more than one
411 model is available, selection favors a conformist bias

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

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

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

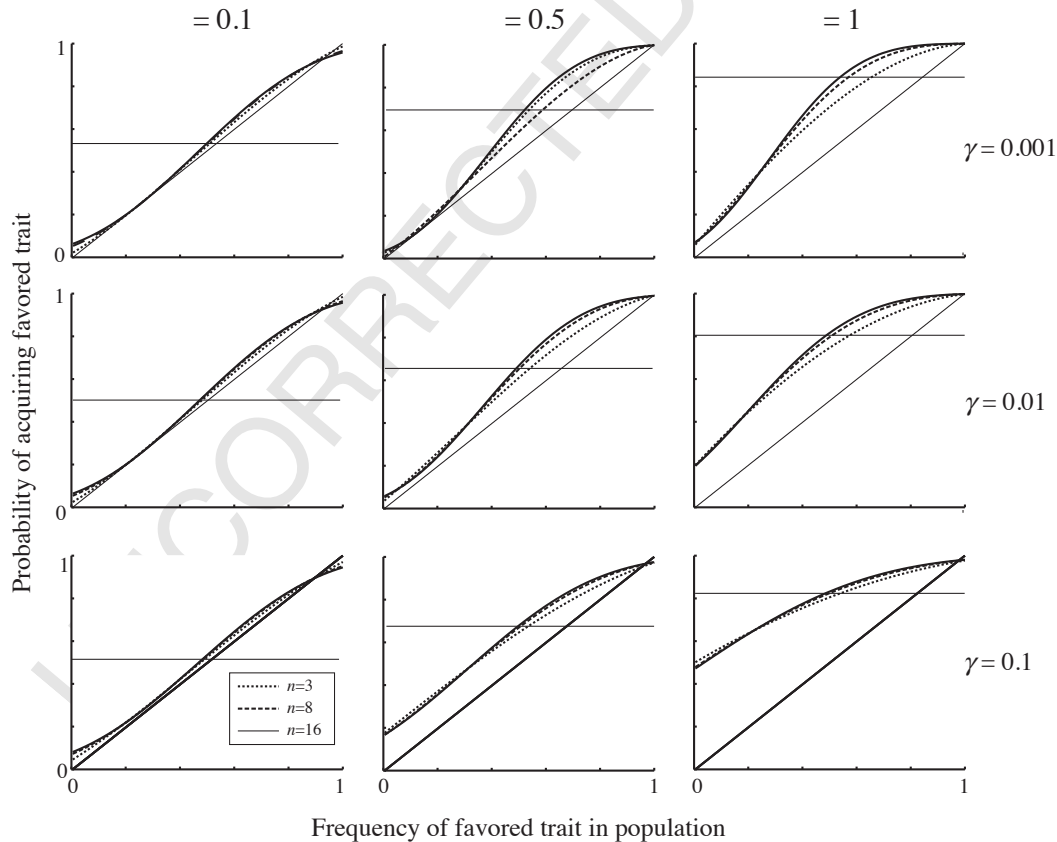


Fig. 2. Probability of acquiring the favored trait as a function of the frequency of the trait in the population, p , at the equilibrium steady state when individuals observe $n=3$, $n=8$, or $n=16$ models for three levels of environmental cue quality (μ) and three different rates of environmental change (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 trait 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 trait is equivalent to the frequency of the favored trait in the population during the next generation (p').

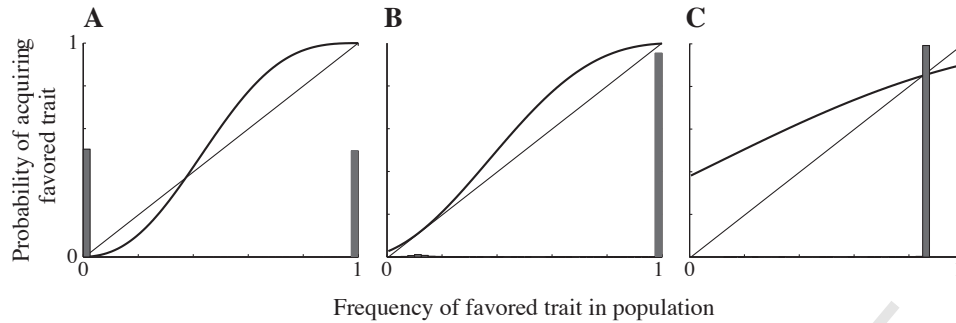


Fig. 3. Probability of acquiring the favored trait as a function of the frequency of the trait, p , for values of g that are lower (A, $g=1$) or higher (C, $g=5$) than the equilibrium value (B, $g=1.66$). The superimposed histograms show the proportion of time steps spent by the population at the various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

453 frequency of the favored trait always increases through time
 454 (unless the environment shifts). The rate of increase is very
 455 slow when it is rare; nonetheless, most of the time, the
 456 favored trait is at very high frequency ($p = .95-1.0$), and
 Q5457 only a small amount of time at frequencies around 0.1.
 458 Lower g values result in individuals that weigh social
 459 information more heavily, and as a result, the favored trait
 460 cannot increase after an environmental change, as is the
 461 case in Fig. 3A. When this happens, the population spends
 462 half the time with the favored trait as the most frequent one
 463 and half the time with the favored trait almost absent. On
 464 the other hand, g values higher than the equilibrium value
 465 cause an individual to weigh environmental information
 466 more heavily. As a result, the favored trait increases more
 467 rapidly after an environmental change, but it also results in
 468 a lower stable equilibrium of p (Fig. 3C). The equilibrium
 469 learning rule is the one that allows the population to
 470 maximize the time spent at high p values (i.e., high
 471 frequencies of the favored trait).

472 Our results suggest that because of this trade-off, selection
 473 is maximizing expected fitness when $n=1$, but bet hedging
 474 when $n>1$. If selection shapes g so as to maximize expected
 Q6475 fitness, the ESS value of g should equal the value given by
 476 Eq. (4) calculated using the observed steady-state frequency
 477 of the favored trait p . If the ESS value of g is larger than the
 478 predicted value, individuals are not maximizing expected
 479 fitness, but instead are relying more on the environmental
 480 cue. Since the environmental cue is noisy, this leads to bet
 481 hedging—increasing the chance of picking the wrong trait so
 482 as to increase the chance of the lineage having some
 483 individuals carrying the best variant after an environmental
 484 switch. The patterns produced for ESS values of g are
 485 consistent with this interpretation. Substituting the actual
 486 steady-state values of the frequency of the favored trait into
 487 Eq. (4) yields a predicted value of g assuming that selection
 488 is maximizing expected fitness. (Fig. 4) plots the difference
 489 between the observed ESS values of g/n and the value of g/n
 490 predicted assuming expected fitness maximization. When
 491 $n=1$, the predicted and computed values are the same on
 492 average, suggesting that selection is maximizing expected

493 fitness. However, when $n>1$, the ESS values of g/n are
 494 systematically larger than the predicted values. In general,
 495 selection maximizes geometric mean fitness in a temporally
 496 varying environment, but in our simulations, the difference
 497 in g values between competing alleles is very small, so
 498 selection is weak, and therefore, geometric mean fitness is
 499 approximately equal to expected fitness. When $n>1$,
 500 however, even small differences in g can lead to big
 501 differences in geometric mean fitness. As the number of
 502 models observed increases, conformism potentially becomes
 503 more costly as it can trap a population at the maladaptive trait
 504 (low values of p). Such deleterious effect of conformism on
 505 geometric mean fitness does not exist when $n=1$, as there can
 506 be no conformism with one model. As a result, we believe
 507 that selection adjusts g so as to create the highest long-term
 508 frequency of the favored trait consistent with not being
 509 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.

510 3.3. Social learning increases average fitness

511 In the simplest models of the evolution of cultural
 512 transmission, social learning evolves but does not increase
 513 average fitness, a phenomenon sometimes labeled “Rogers
 514 Paradox” (Rogers, 1988). Average fitness increases if
 515 models and learners are related (Lehmann et al., 2010;
 516 Rendell, Boyd et al., 2010; Rendell, Fogarty et al., 2010)
 517 or if social learning allows for the selective use of
 518 environmental information (i.e., use environmental infor-
 Q7519 mation when it is accurate; otherwise, imitate) (1995;
 520 Enquist et al., 2007).

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

quality of the environmental cue is poor (when μ is low). 528
 The fitness advantage associated with social learning is 529
 greater in a stable environment (low γ), as higher γ increases 530
 the likelihood that the social information in a population is 531
 outdated and inaccurate. 532

533 Interestingly, increasing the number of models some-
 534 times decreases average fitness. Fig. 5 shows how the 534
 number of individuals observed (n) impacts the relative 535
 average fitness advantage of social learning. When the 536
 environment is unstable ($\gamma=0.1$, Fig. 5A), expected fitness 537
 increases with n . In more stable environments ($\gamma=0.01$), the 538
 fitness advantage of social learning as a function of n is 539
 concave with a peak at $n=3$ (Fig. 5B). This suggests that 540
 when the environment is stable and environmental cue is 541
 noisy, there are an optimal number of models, and paying 542
 attention to more models decreases expected fitness. We 543
 believe that this is an artifact that results from the fact that, 544
 in the current model, learners have a very simple prior 545
 belief about the long run average frequency of the favored 546

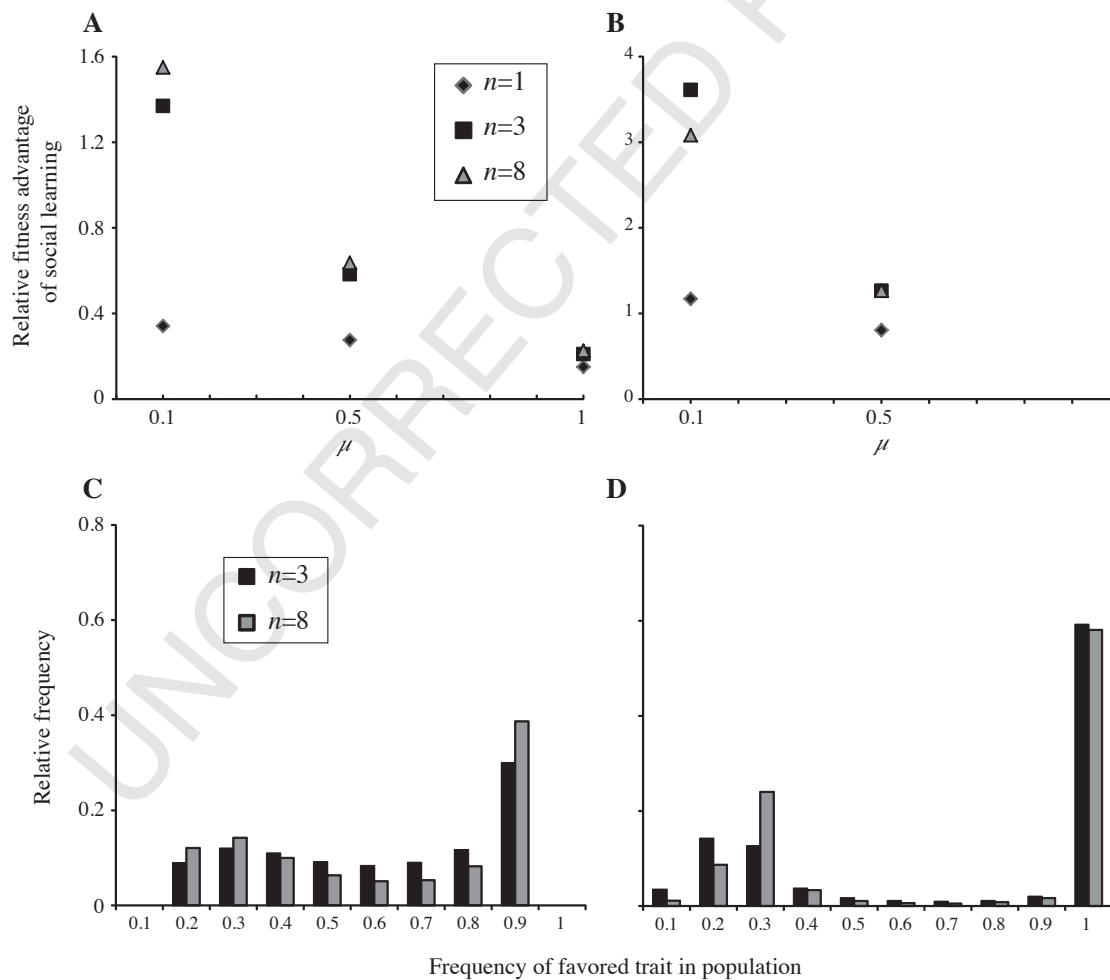


Fig. 5. Fitness of the evolutionary equilibrium learning rules relative to the expected fitness of pure individual learners for different values of n when (A) $\gamma=0.1$ and (B) $\gamma=0.01$. The histograms (C) and (D) show the relative frequency of time spent by the population at different values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

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

558 3.4. Imitation error or migration increases reliance on 559 social information when more than one model is available

560 In previous work, it has been argued that both spatial
 561 variation in fitness (McElreath et al. (in press)) and
 562 inaccurate imitation (Henrich & Boyd, 2002) lead favor
 Q8563 an increased the reliance on conformism. Here we describe
 564 a modification of the model to allow for spatial variation or
 565 imitation error. Assume with probability $1-\epsilon$ individuals
 566 accurately observe the behavior of a model, but with
 567 probability ϵ , the learner believes that the model to has
 568 behavior 1 50% of the time, and behavior 2 1 50% of the
 Q9569 time. This can be thought of as due to a failure to
 570 accurately perceive the model's behavior or due to an
 571 influx of migrants from a large population in which $p = .5$.
 572 This would be the case in an infinite island model in which
 573 the environments in different islands fluctuated indepen-
 574 dently. Under these conditions, a migration rate of 0.1
 575 would mean that, 10% of the time, a learner would learn
 576 from a migrant, who is expected to have behavior 1 50% of
 577 the time. Thus, the population-level consequence of a
 578 migration rate of 0.1 is equivalent to a rate of observation
 579 failure of 0.1.

580 Adding imitation error or migration ($\epsilon=0.1$) decreases
 581 the accuracy of social information—the observed behavior
 582 of others is a poorer predictor of the environmental state
 Q1083 (see the SOM). When learners can only observe one
 584 model ($n=1$), this leads to an increased reliance on
 585 environmental information. Surprisingly, however, when
 586 individuals have access to more than one model ($n>1$),
 587 adding imitation error or migration actually *decreases*
 588 the equilibrium value of g and thus increases the reliance on
 589 social cues (see Table 1 in the SOM, available on the
 590 journal's website at www.ehbonline.org). This interesting
 591 result reflects the fact that increasing g also increases the
 592 degree of conformism, and this allows learners to
 593 circumvent the deleterious effect of imitation error when
 594 $p>.5$. By observing more than one individual, learners
 595 can conform to the majority of models, which is in
 596 general still accurately perceived since the magnitude of
 597 imitation errors is not sufficiently large to make the
 598 unfavored trait the dominant one when the favored trait is
 599 dominant and vice versa. However, average population
 600 fitness decreases when imitation errors are added.

4. Discussion

601

The model of the evolution of social learning analyzed 602
 above is based on two assumptions. The first one is that 603
 natural selection shapes learning mechanisms so that 604
 alternative cues about the environment are weighted 605
 optimally according to Bayes' law so that reliance on a cue 606
 depends on the extent to which that cue predicts fitness- 607
 relevant contingencies. The second assumption is that the 608
 predictive value of other individuals' observed behavior 609
 depends on the past history of social learning and 610
 environmental change. This means that determining the 611
 evolutionarily stable social learning psychology requires 612
 modeling both the coevolution of the pool of culturally 613
 transmitted information and the genes that determine how 614
 individual psychology processes this information. 615

Taking this approach, we have shown that selection 616
 favors a psychology that causes individuals to rely heavily 617
 on imitating the behavior of others when environmental 618
 cues are not very informative and environments change 619
 slowly. We also show that cultural transmission errors and 620
 migration from other environments increase reliance on 621
 imitation when learners can observe the behavior of more 622
 than one cultural model. These results replicate previous 623
 work on the topic (Boyd & Richerson, 1985, 1987; Henrich 624
 & Boyd, 1998; McElreath et al. in press) without making 625
 the assumption that social learning and individual learning 626
 are distinct processes. This suggests that the qualitative 627
 conclusions drawn from these studies and this paper are 628
 robust and do not hinge on the specificity of any particular 629
 learning mechanism. 630

We have also found that a conformist bias is likely to be a 631
 universal aspect of social learning consistent with some 632
 previous work (Boyd & Richerson, 1982, 1985; Henrich & 633
 Boyd, 1998). Understanding the conditions under which 634
 conformist-biased social learning evolves is important 635
 because it could be a human psychological adaptation, and 636
 that can maintain stable between-group cultural differences 637
 which create the conditions for group selection, allowing 638
 large-scale cooperation and altruistic punishment to evolve 639
 (Boyd & Richerson, 1985; Henrich & Boyd, 1998; 640
 Richerson & Boyd, 2005). 641

Recently, there has been debate about the conditions 642
 under which conformism will evolve. Eriksson et al. (2007), 643
 Wakano and Aoki (2007), McElreath et al. (in press), Boyd 644
 and Richerson (1985), Andrés Guzmán and Rodríguez- 645
 Sickert (2007), and Henrich and Boyd (1998) analyzed 646
 models similar to ours in which the environment changes 647
 spatially and temporally between two states. Through 648
 individual learning or cultural transmission, individuals 649
 adopt one of two behaviors, each of which is favored in one 650
 of the environmental states. They show that conformism is 651
 favored over unbiased imitation by natural selection under a 652
 broad range of environmental conditions. Erikson, Enquist, 653
 and Ghirlanda (2007) have criticized these models on the 654 Q11
 ground that they assume that individuals have inherent 655

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

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

706 Bayesian methods are particularly well suited to the study
 707 of the evolution of cultural evolution. We believe that
 708 cultural transmission, as an adaptation, has been designed to
 709 solve an inference problem: given some environmental and
 710 social information, how should an individual infer the state
 711 of its current local environment in order to adopt the

behavior that will maximize its fitness? Learning about the 712
 state of the world, whether through an individual trial-and- 713
 error process or through social imitation, can be conceived of 714
 as a Bayesian process. This study shows that Bayesian tools 715
 can be profitably applied to the study of the evolution of 716
 social learning (Boyd & Richerson, 1995). 717 **Q12**

Supplementary Materials 718

Supplementary data to this article can be found online at 719
 doi:10.1016/j.evolhumbehav.2011.12.007. 720

References 721

- Andrés Guzmán, R., Rodríguez-Sickert, C. et al. (2007). When in Rome, do 722
 as the Romans do: the coevolution of altruistic punishment, conformist 723
 learning, and cooperation. *28*, 112–117. 724
- Benskin, C. M. H., & Mann, N. I., et al. (2002). Social learning directs 725
 feeding preferences in the zebra finch, *Taeniopygia guttata*. *Animal* 726
Behavior, *64*, 823–828. 727
- Boyd, R., & Richerson, P. (1982). Cultural transmission and the evolution of 728
 cooperative behavior. *Human Ecology*, *10*, 325–351. 729
- Boyd, R., & Richerson, P. J. (1985). *Culture and the evolutionary process*. 730
 London: University of Chicago Press. 731
- Boyd, R., & Richerson, P. J. (1987). The evolution of ethnic markers. 732
Cultural Anthropology, *2*, 65–79. 733
- Boyd, R., & Richerson, P. J. (1995). Why does culture increase human 734
 adaptability? *Ethology and Sociobiology*, *16*, 125–143. 735
- Brown, C., & Laland, K. N. (2003). Social learning in fishes: a review. *Fish* 736
and Fisheries, *4*, 280–288. 737
- Cavalli-Sforza, L. L., & Feldman, M. V. (1981). *Cultural transmission and* 738
evolution: a quantitative approach. Princeton: Princeton University 739
 Press. 740
- Enquist, M., & Eriksson, K., et al. (2007). Critical social learning: a solution 741
 to Rogers's Paradox of nonadaptive culture. *American Anthropologist*, 742
109, 727–734. 743
- Eriksson, K., & Enquist, M., et al. (2007). Critical points in current theory of 744
 conformist social learning. *Journal of Evolutionary Psychology*, *5*, 745
 67–87. 746
- Galef, B. G., & Laland, K. N. (2005). Social learning in animals: empirical 747
 studies and theoretical models. *Bioscience*, *55*, 489–499. 748
- Gillespie, J. H. (1977). Natural selection for variances in offspring numbers: 749
 a new evolutionary principle. *The American Naturalist*, *111*, 750
 1010–1014. 751
- Henrich, J., & Boyd, R. (1998). The evolution of conformist transmission 752
 and the emergence of between-group differences. *Evolution and Human* 753
Behavior, *19*, 215–241. 754
- Henrich, J., & Boyd, R. (2002). On modeling cognition and culture: why 755
 replicators are not necessary for cultural evolution. *Journal of cognition* 756
and culture, *2*, 87–112. 757
- Heyes, C. M. (1994). Social learning in animals: categories and 758
 mechanisms. *Biological Review*, *69*, 207–232. 759
- Kalish, M. L., & Griffiths, T. L., et al (2007). Iterated learning: 760
 intergenerational knowledge transmission reveals inductive biases. 761
Psychonomic Bulletin & Review, *14*, 288–294. 762
- Kameda, T., & Nakanishi, D. (2003). Does social/cultural learning increase 763
 human adaptability? Rogers's question revisited. *Evolution and Human* 764
Behavior, *24*, 242–260. 765
- Laland, K. N. (2004). Social learning strategies. *Learning & Behavior*, *32*, 766
 4–14. 767
- Leadbeater, E., & Chittka, L. (2007). Social learning in insects—from 768
 miniature brains to consensus building. *Current Biology*, *17*, 769
 R703–R713. 770

- 771 Lefebvre, L. (2000). Feeding innovations and their cultural transmission in
772 bird populations. In C. Heyes, & L. Huber (Eds.), *The evolution of*
773 *cognition* (pp. 311–328). London: MIT Press.
- 774 Lehmann, L., & Feldman, M. W., et al. (2010). Cumulative cultural
775 dynamics and the coevolution of cultural innovation and transmission:
776 an ESS model for panmictic and structured populations. *Journal of*
777 *Evolutionary Biology*, 23, 2356–2369.
- 778 McElreath, R., Fasolo, B., et al. (in press). The evolutionary rationality of
social learning In R. Hertwig and U. Hoffrage, MIT Press.
- Q1379 McElreath, R., & Strimling, P. (2008). When natural selection favors
780 imitation of parents. *Current Anthropology*, 49, 307–316.
- 781 Mesoudi, A. (2011). *Cultural evolution*. Chicago: University of Chicago
782 Press.
- 783 A. Perfors, J.B. Tenenbaum, et al., A tutorial introduction to Bayesian
784 models of cognitive development. *Cognition*. In press, corrected proof.
- Q1485 Perry, S. E., & Manson, J. H. (2003). Traditions in monkeys. *Evolutionary*
786 *Anthropology: Issues, News, and Reviews*, 12, 71–81.
- 787 Plotkin, H. C. (1988). Learning and evolution. In H. C. Plotkin, Ed. *The role*
788 *of behavior in evolution* (pp. 133–164). Cambridge, MA: MIT Press.
- Rendell, L., & Boyd, R., et al. (2010). Why copy others? Insights from the
790 social learning strategies tournament. *Science*, 328, 208–213. 791
- Rendell, L., & Fogarty, L., et al. (2010). Rogers' paradox recast and resolved:
792 population structure and the evolution of social learning strategies. 793
Evolution, 64, 534–548. 794
- Richerson, P. J., & Boyd, R. (2005). *Not by genes alone: how culture*
795 *transformed human evolution*. Chicago: University of Chicago Press; 796
2005. 797
- Rogers, A. R. (1988). Does biology constrain culture? *American*
798 *Anthropologist*, 90, 819–831. 799
- Tenenbaum, J. B., & Griffiths, T. L., et al. (2006). Theory-based Bayesian
800 models of inductive learning and reasoning. *Trends in Cognitive*
801 *Sciences*, 10, 309–318. 802
- Wakano, J. Y., & Aoki, K. (2007). Do social learning and conformist bias
803 coevolve? Henrich and Boyd revisited. *Theoretical Population Biology*,
804 72, 504–512. 805
- Whiten, A., & Hinde, R. A., et al. (2001). Discussion meeting issue 'culture
806 evolves'. *Philosophical Transactions of the Royal Society B-Biological*
807 *Sciences*, 366. 808
809

SOM EHB2.