

THE COEVOLUTION OF SOCIAL LEARNING AND SENSITIVITY TO CHANGING ENVIRONMENTS

RICHARD MCELREATH*

ABSTRACT. There is widespread agreement that social and individual learning are adaptations to varying environments. However, existing theory assumes that organisms cannot detect changes in the environment and instead adapt to averages. This paper develops the first analytical model that allows for the simultaneous coevolution of socially learned traditions, reliance on social learning, and signal detection for environmental change. There are numerous conditions under which detection can be stable once common but cannot invade the population when rare. When signal detection is maintained by selection, it always leads to pure separating equilibria at which organisms always learn individually when they believe the environment has recently changed and otherwise always learn socially. Detection can increase mean fitness at equilibrium, but it may also reduce it.

1. INTRODUCTION

1
2 This paper asks when natural selection will favor an organism's paying a
3 cost to detect changes in the environment, provided that cues of environ-
4 mental change adjust use of individual and social learning. I use formal
5 modeling to address this question. But the motivation for the paper is re-
6 ally empirical, meant to address a gap between the structure of the theory
7 and how it is interpreted in light of data. When I was a new assistant pro-
8 fessor, I set out with colleagues Peter J. Richerson, Mark Lubell and several
9 industrious PhD students to follow Kameda and Nakanishi (2002) and de-
10 velop an experimental program for studying the adaptive design of social
11 learning in humans (McElreath *et al.* 2005, 2008, Efferson *et al.* 2007, 2008).
12 The goal was to evaluate the predictions of theory, by using different experi-
13 mental treatments to simulate differences in theoretical parameters, such as
14 rate of environmental change, that lead to changing predictions for reliance
15 on social learning. A number of other laboratories have also explored the

Date: June 6, 2012.

*Department of Anthropology and Center for Population Biology, UC Davis, One Shields Avenue, Davis CA 95616. mcelreath@ucdavis.edu.

16 same sorts of questions, and results confirm the qualitative agreement be-
17 tween what the models predict and how real people (or at least, real univer-
18 sity students) change their social learning strategies in response to changes in
19 aspects of the social and physical environment (Morgan *et al.* 2012, provide
20 a recent and clear example).

21 But a number of issues with interpreting the experimental results have pre-
22 occupied us. For example, laboratory social learning experiments rarely cre-
23 ate the differences in experience that both natural communities and theoret-
24 ical models always possess. Instead, a group of naive participants are asked
25 to learn from one another, and their behavior is predicted using models that
26 assume a different life history, in which naive individuals always coexist with
27 experienced individuals. As a result, those rare experiments that do estab-
28 lished “overlapping” generations may show much more powerful and realis-
29 tic social learning effects (Baum *et al.* 2004, Jacobs and Campbell 1961).

30 Another concern, and the one that occupies the remainder of this paper,
31 has been that the formal theory itself does not allow for the kind of savvy
32 attention to contextual variables that our design and interpretation of the
33 experiments assumes. We have been expecting savvy context-sensitive de-
34 ployment of individual and social learning strategies, based upon the inter-
35 pretation of formal models in which learning strategies respond to context
36 only over evolutionary time (reviewed in McElreath *et al.* in press). Even the
37 interpretation of my own fieldwork has used this slight of hand (McElreath
38 2004). There are a few cases in which models have allowed for contingent
39 strategy use (Henrich and Boyd 1998, Boyd and Richerson 1996, Enquist
40 *et al.* 2007, McElreath *et al.* 2008). However, for the most part the literature
41 has focused on how evolution, rather than individuals, could strategically
42 adjust learning.

43 This focus has made it difficult to really say what theory predicts. It makes
44 sense to view the evolution of contingent social learning as a special case of
45 the general theory of phenotypic plasticity. Social learning is itself a form
46 a phenotypic plasticity, and the evolution of plastic use of it is a kind of
47 meta-plasticity. We might wonder when such meta-plasticity might evolve,
48 because the general evolutionary ecology literature has long confirmed that
49 phenotypic plasticity is not always favored by natural selection (Levins 1968,
50 DeWitt *et al.* 1998). The general literature on the evolution of phenotypic
51 plasticity is too vast to review here, but it is worth noting that selection may
52 not favor an organism’s adjusting phenotype in response to cues (e.g. Co-
53 hen 1967, Getty 1996, Tufto 2000), and non-contingent strategies can be
54 favored even when environmental cues are reliable (McNamara and Dall
55 2011). These results have been recently generalized to a wide range of biolog-
56 ical phenomena (Altenberg 2012, in press). Some of my own previous theory

57 on the evolution of social learning has turned out to illustrate it (McElreath
58 and Strimling 2008, as cited in McNamara and Dall 2011). On the other
59 hand, in a recent high-profile simulation tournament of the evolution of so-
60 cial learning strategies, one of the most successful strategies regulated social
61 learning by using the time since behavior was learned, in combination with
62 an estimate of how quickly payoffs change over time (Rendell *et al.* 2010).

63 In light of these results, it's worth wondering when we should expect peo-
64 ple and other organisms to pay attention to cues that regulate learning strat-
65 egy. We might begin by reconsidering formal models of gene-culture co-
66 evolution in the presence of sensitivity to changing environments. When
67 will natural selection favor using cues of spatial or temporal environmental
68 change to regulate mode of learning? Ongoing debates about the adaptive-
69 ness of strategies such as conformist transmission, which has long been cen-
70 tral to the gene-culture coevolution literature (Boyd and Richerson 1985),
71 may depend upon understanding selection for such sensitivity (McElreath
72 *et al.* in press, Nakahashi *et al.* in press). And as the planet warms and is
73 otherwise rapidly altered by human activity, predicting and understanding
74 species' responses will partially depend upon our ability to make sense of the
75 design of environmental sensitivity (Sih *et al.* 2011).

76 The rest of this paper develops a first model that directly addresses the
77 question: *When will natural selection favor attention to cues of temporally*
78 *changing environments in order to regulate reliance on individual and social*
79 *learning?* I use a common gene-culture or dual-inheritance modeling frame-
80 work (Cavalli-Sforza and Feldman 1981, Boyd and Richerson 1985). I add to
81 this structure another heritable component of plasticity that invests in detect-
82 ing temporal changes in the environment. The organism can use different
83 learning strategies depending upon whether or not it believes the environ-
84 ment has recently changed. Using a signal detection framework, like a num-
85 ber of previous theoretical studies of phenotypic plasticity (e.g. Getty 1996),
86 I show that gene-culture coevolution may lead to substantial investments in
87 detecting change, but that such investment is not always favored. Indeed, the
88 range of conditions that can stabilize sensitivity to changing environments is
89 always larger than the range that will allow it to invade the population. But
90 whenever detection does evolve, it leads to a perfect separating equilibrium
91 at which the organism always learns individually, when it believes the envi-
92 ronment has recently changed, and otherwise always learns socially, when it
93 believes the environment has not recently changed. The result is that much
94 more social learning is observed, once detection evolves. Despite the increase
95 in the amount of social learning, the expected population growth rate may
96 nevertheless increase in the presence of detection, due to adaptive allocation
97 of individual learning to time periods in which it is needed most. I close the

98 paper by considering limits of the model, avenues for future work, and the
99 impact of these results on the interpretation of empirical evidence.

100 2. MODEL ASSUMPTIONS

101 For comparability to existing theory, I use a traditional discrete genera-
102 tion, infinite population framework to construct the model (Cavalli-Sforza
103 and Feldman 1981, Boyd and Richerson 1985, Rogers 1988). Many models
104 address a similar core problem. The adaptive challenge for the organism is
105 to acquire optimal behavior in a temporally varying environment. Since the
106 optimal behavior changes over time, always learning socially is never evo-
107 lutionarily stable. But similarly, since asocial learning is more costly, unless
108 optimal behavior changes very quickly, some social learning is usually fa-
109 vored. The rate of environmental change and the cost of asocial learning
110 govern the evolutionarily stable mix of social and asocial learning. Because
111 of geometric mean fitness effects and bet hedging, natural selection tends
112 to favor mixed learning strategies over pure ones (McElreath *et al.* in press,
113 Perreault *et al.* in press).

114 In this paper, I use continuous strategy spaces, allowing individual geno-
115 types to code for probabilities of individual and social learning in different
116 contexts. I keep as much as possible about the core model the same as other
117 papers. The traditional framework has several drawbacks, which I explore
118 in the discussion. However, I wish to begin by changing as little as possi-
119 ble about existing theory, in order to understand the consequences of allow-
120 ing sensitivity to changing environments to regulate social learning. I intro-
121 duce into the basic model the ability for an organism to detect environmental
122 change and use different probabilities of social learning depending upon its
123 inference. I develop a weak selection approximation to the geometric mean
124 fitness of a mutant, which allows me to define the evolutionary dynamics of
125 detection. The rest of this section defines the model in detail.

126 **2.1. Population and life cycle.** Suppose a large well-mixed population of
127 semelparous organisms that are capable of both individual and social learn-
128 ing. The environment the organisms inhabit is everywhere the same, but may
129 change from one generation to the next. Let u be the chance of the environ-
130 ment changing in any given generation. The current state of the environment
131 prescribes a unique behavior that results in an increase in expected reproduc-
132 tion (“fitness”) b . All other behavior results in no change in fitness. When
133 the environment changes, it changes to a new state it has never had before,
134 and all previous behavior is rendered non-optimal.

135 **2.2. Heritable strategies.** Behavior is always acquired via learning. But learn-
136 ing strategy is a heritable trait that specifies the probability of using individual

137 learning, instead of social learning. Employing individual learning means
138 that the organism pays a fitness cost bk (a proportion k of the maximum
139 gain b) for a chance s of learning optimal behavior. Social learning means
140 that an individual pays no up-front learning cost relative to individual learn-
141 ing, but instead copies a random member of the previous generation. While
142 cheaper, social learning may or may not yield currently optimal behavior,
143 and so it may ultimately be more expensive than individual learning, espe-
144 cially just after a change in the environment.

145 The adaptive challenge the model explores is how individuals regulate their
146 learning strategy, based upon information that the environment has recently
147 changed. Let p_s be the heritable probability of deploying individual learning
148 when an individual believes the environment has been *stable*, since the last
149 generation. Let p_c be the probability of deploying individual learning when
150 the individual believes the environment has *changed* since the last generation.

151 **2.3. Signal detection.** Individuals acquire beliefs about the state of the envi-
152 ronment via investment in detecting signals of recent change. These signals
153 may be anything from changes in perceived efficacy of a technology or tech-
154 nique to appreciation of others' opinions on whether or not the environment
155 has changed. I comment more on the nature of such signals in the discussion.
156 The crucial limiting assumption in this model will be that it is individual in-
157 vestments that affect belief formation. Let d be the probability of correctly
158 detecting a change in the environment. This is an individual heritable char-
159 acter, with population mean d^* . Let $f(d)$ be a function that determines the
160 probability of a false positive, of thinking the environment changed when it
161 did not. The population mean rate of false positives is $f(d^*) = f^*$.

162 I leave this function undefined for now. However, there are several limit-
163 ing assumptions we can make about the shape of this function, before defin-
164 ing it, and these assumptions will be sufficient to prove the invasion and sta-
165 bility criteria for the model. The general shape of the function $f(d)$ comes
166 from analogy to a Receiver Operating Characteristic (Green and Swets 1966).
167 A Receiver Operating Characteristic (ROC) describes the tradeoff between
168 accuracy and error in a classification task. As the ability of a signal or test
169 to detect true cases rises, so too does the rate of false positives. As a result,
170 optimal detection in real classification tasks almost always means accepting
171 some false-negatives as well as false-positives.

172 Readers familiar with the signal detection literature will recognize d as the
173 sensitivity and $f(d)$ as one-minus-specificity, the Type II error rate. The ex-
174 act shape of the tradeoff between detection and false alarms depends upon
175 the details of each case, but the general nature of this tradeoff is nearly uni-
176 versal in signal detection. Figure 1 illustrates the general shape $f(d)$ must
177 take. First, I restrict $f(d)$ to continuous, differentiable functions. Second,




FIGURE 1. Relationship between true detection rate d to false positives $f(d)$. Increasing investments in d on the horizontal axis lead to increases in the rate of false positives, $f(d)$. The dashed line on the diagonal represents $f(d) = d$, where the signal is useless because the rate of true positives equals the rate of false positives. Below this line, the signal provides information that allows individuals to detect true changes in the environment. Example curves in this figure are the function $f(d) = ad/(1 + a - d)$. The solid curve is for $a = 0.1$. The dashed curve is for $a = 1$. Larger values of a indicate higher rates of false positives for any given rate of true positives.

178 realistic signal detection problems have a few recurring features. The only
 179 reliable way to detect all true cases is to always assume that the event has oc-
 180 curred. This implies that $f(1) = 1$ —if the detection rate is 100%, then the
 181 false positive rate is also 100%. Likewise, the only way to miss every true
 182 case is to assume the event never happens, $f(0) = 0$. Third, I assume that
 183 $f(d) \leq d$, the rate of false positives is everywhere less than the rate of true
 184 detection, unless $d = 1$. Finally, the previous assumptions imply that the
 185 rate of change in false positives is everywhere positive or zero, $f'(d) \geq 0$,
 186 and that the acceleration of false positives is strictly positive, $f''(d) > 0$. It is
 187 also necessary that $f'(0) < 1$, as a consequence of assuming $f(d) \leq d$.

188 Although I will prove most of the interesting features of this model for any
 189 $f(d)$ that fits the restrictions above, in order to illustrate the dynamics of the
 190 model, I will later need a particular function $f(d)$. Specifically, I will use a

191 flexible hyperbolic function for all numerical examples in this paper:

$$f(d) = \frac{ad}{1 + a - d}, \quad (1)$$

192 where $a > 0$ is a parameter that determines how quickly $f(d)$ increases as
 193 d increases. This function ensures that when $d = 1$ (detection is perfectly
 194 accurate), that $f(1) = 1$ also (false positives always occur). At the limit $a \rightarrow$
 195 0 , $f(d) = 0$, the signal is perfect. But at the limit $a \rightarrow \infty$, $f(d) = d$ and the
 196 signal is useless, because behaving according to the signal is just like guessing.
 197 Figure 1 plots this function for $a = 0.1$ and $a = 1$.

198 **2.4. The cost of signal detection.** Attempting to detect environmental change
 199 carries a fixed fitness cost, $b\ell d$, where $\ell > 0$ is a new parameter that governs
 200 the marginal cost of signal detection. Individuals who are increasingly sensi-
 201 tive to environmental change pay an increasing fitness cost. This assumption
 202 allows for a wide range of different mechanistic hypotheses. If we suspect
 203 that information about environmental change is quite cheap to acquire and
 204 process, then ℓ can be made to be close to zero. If we suspect instead that
 205 such information is costly to acquire or process, then ℓ will be large.

206 **2.5. Fitness at time t .** With the assumptions above, we can write a gener-
 207 al fitness function for a mutant individual with individual learning prob-
 208 abilities p_s and p_c and detection rate d in a population in which everyone
 209 else has probabilities p_s^*, p_c^*, d^* . Let w_0 be baseline fitness accrued through
 210 other activities. Let t be the number of generations since the last change in
 211 the environment. For notational simplicity, I define $f \equiv f(d)$ and $f^* \equiv$
 212 $f(d^*)$. When the environment has changed since the previous generation
 213 completed their learning, $t = 0$, detection of true change affects fitness. Then
 214 the expected fitness of the mutant is:

$$w_{t=0}(p_s, p_c, d, p_s^*, p_c^*, d^*) = w_0 + dp_c b(s - k) + (1 - d)p_s b(s - k) - b\ell d.$$

215 The fitness of this mutant at a time $t > 0$ generations since the last change in
 216 the environment is:

$$w_{t>0}(p_s, p_c, d, p_s^*, p_c^*, d^*) = w_0 + (1 - f)(p_s b(s - k) + (1 - p_s)q_t b) \\ + f(p_c b(s - k) + (1 - p_c)q_t b) - b\ell d,$$

217 where $q_t = q_t(p_s^*, p_c^*, d^*)$ is a function that yields the probability of acquir-
 218 ing optimal behavior via social learning, t generations after a change in the
 219 environment. I derive q_t in the next section.

220 All together the fitness of a mutant t generations after the most recent
 221 change in the environment is given by:

$$w_t(p_s, p_c, d, p_s^*, p_c^*, d^*) \equiv w_t = \quad (2)$$

$$w_0 + \begin{cases} b(s-k)E_c - bld & \text{if } t = 0 \\ b(E_s(s-k) + (1-E_s)q_t) - bld & \text{if } t > 0 \end{cases},$$

222 where $E_c = dp_c + (1-d)p_s$ is a mutant's expected amount of individual
 223 learning, just after a change in the environment, and $E_s = fp_c + (1-f)p_s$ is
 224 the mutant's expected amount of individual learning when the environment
 225 has not recently changed.

226 **2.6. Quality of social information, q_t .** The next step in computing the growth
 227 rate of the mutant strategy is to compute q_t , the probability of acquiring adap-
 228 tive behavior via social learning, t generations after the most recent change in
 229 the environment. The problem is to define the recurrence process by which
 230 adaptive behavior accumulates in the population. Just after a change in the
 231 environment ($t = 0$), there is no chance of acquiring adaptive behavior via
 232 social learning, because all behavior that was learned in previous generations
 233 is now non-optimal. Every generation that the environment remains stable,
 234 adaptive behavior is pumped into the population via the action of individual
 235 learning.

236 In the Supporting Information, I use the logic above to derive the ex-
 237 plicit function for q_t , the probability of acquiring adaptive behavior via social
 238 learning, t generations after a change in the environment:

$$q_t = \begin{cases} 0 & \text{if } t = 0 \\ s(1 - (1 - E_c^*)(1 - E_s^*)^{t-1}) & \text{if } t > 0 \end{cases}, \quad (3)$$

239 where $E_c^* = d^*p_c^* + (1-d^*)p_s^*$ is the average amount of individual learning in
 240 the population, just after a change in the environment, and $E_s^* = f^*p_c^* + (1-$
 241 $f^*)p_s^*$ is the average amount of individual learning when the environment
 242 has not recently changed. I use this function in the next section to estimate
 243 the growth rate of the mutant.

244 **2.7. Long run expected growth rate.** The probability that the mutant will
 245 increase in frequency depends upon the stochastic nature of the environ-
 246 ment. To compute the required expression, we note that selection in time
 247 varying environments, at least with simple life histories such as these, will
 248 maximize the geometric mean fitness, not the arithmetic mean fitness. For a
 249 particularly clear explanation of this fact, see Cohen (1966). I label the geo-
 250 metric mean fitness of the mutant $r(p_s, p_c, d, p_s^*, p_c^*, d^*)$ and work with its

251 natural logarithm. This expression is defined as:

$$\log r(p_s, p_c, d, p_s^*, p_c^*, d^*) \equiv \log(r) = \sum_{t=0}^{\infty} \Pr(t) \log(w_t).$$

252 This is just the natural logarithm of the geometric mean fitness of the invad-
253 ing mutant.

254 The only part of this puzzle still missing is a function defining $\Pr(t)$, the
255 chance the environmental takes on the state t in any given generation. This
256 is given by $\Pr(t) = u(1-u)^t$. If the environment just changed, then $t = 0$,
257 and this happens with probability u , by the definition of u . In order to reach
258 $t = 1$, the environment has to remain stable for one generation. The chance
259 of this is $u(1-u)$. For $t = 2$, the chance must be $u(1-u)^2$, because a sequence
260 of two generations without a change is necessary. A similar derivation of this
261 geometric relationship appears in Rogers (1988).

262 **2.8. Weak selection approximation.** The expression $\log(r)$ above is incon-
263 venient for analysis. There is no known method for closing this kind of in-
264 finite series, in which the index variable t is an exponent both inside and
265 outside of the logarithm. To make progress, I use the customary tactic. I
266 construct a weak selection approximation by using a Taylor series expansion
267 of $\log(r)$ around the point $b = 0$ and keeping the linear term only. This
268 provides an approximation of the model for $b^2 \approx 0$, corresponding to the
269 assumption that selection is weak:

$$\log(r) \approx \log(w_0) + \frac{b}{w_0} (I + S) - bld. \quad (4)$$

270 S is a term summarizing the fitness benefits of social learning, and I is a term
271 summarizing the fitness benefits of individual learning. These terms are:

$$I = (s - k)(p_s + (p_c - p_s)(ud + (1 - u)f)), \quad (5)$$

$$S = (1 - u) \overbrace{(1 - p_s - (p_c - p_s)f)}^L \overbrace{\left(s \frac{p_s^* + (p_c^* - p_s^*)(ud^* + (1 - u)f^*)}{1 - (1 - u)(1 - p_s^* - (p_c^* - p_s^*)f^*)} \right)}^Q. \quad (6)$$

272 Some sense can be made of these expressions, before analyzing the dynam-
273 ics. Consider the expression for I . It is proportional to $s - k$, the propor-
274 tion of fitness benefits that remain after subtracting the costs of individual
275 learning. The rest of the expression merely quantifies the mutant's rate of in-
276 dividual learning, taking into account signals of environmental change and
277 the different rates of learning they create. Note that the common-type trait
278 values p_s^*, p_c^*, d^* do not appear in the expression for I .

279 Social learning, expression S , however does depend upon common-type
 280 strategy. Social learning only pays when the environment has been stable
 281 for at least one generation, and so the entire expression is proportional to
 282 $1 - u$. The term labeled L is the rate of social learning for the mutant, when
 283 the environment is stable. The term labeled Q quantifies the expected qual-
 284 ity of social information. It is exactly the expected probability of acquiring
 285 adaptive behavior via social learning, conditioned on the environment being
 286 stable ($t > 0$). It depends upon the common type phenotypes p_s^* , p_c^* and d^* ,
 287 because the common type creates the cultural environment that the mutant
 288 experiences. Note that the numerator of this term is just the common-type
 289 rate of individual learning, which is the rate at which new adaptive behavior
 290 enters the population. This input of adaptive behavior is discounted by the
 291 denominator, which is one minus the probability of social learning, given
 292 that the environment is stable. As the amount of social learning increases,
 293 the denominator gets smaller, making any inputs from individual learning
 294 accumulate more. So the denominator in total can be thought of as a cultural
 295 turnover rate. When it is small, because social learning is common and the
 296 environment is relatively stable, the entire value of Q is increased through
 297 accumulation of past innovations. When the denominator is instead small,
 298 because either social learning is rare or the environment is relatively unsta-
 299 ble, then Q is reduced.

300

3. ANALYSIS

301 To analyze the model, I use a tactic common in evolutionary ecology and
 302 evolutionary game theory. If mutants are rare and phenotypically very close
 303 to the common type, then the change in each trait is proportional to the rate
 304 of change in mutant fitness:

$$\begin{aligned} \Delta p_s^* \propto \partial \log(r)_{p_s} &\equiv \left. \frac{\partial \log(r)}{\partial p_s} \right|_{p_s=p_s^*, p_c=p_c^*, d=d^*}, \\ \Delta p_c^* \propto \partial \log(r)_{p_c} &\equiv \left. \frac{\partial \log(r)}{\partial p_c} \right|_{p_s=p_s^*, p_c=p_c^*, d=d^*}, \\ \Delta d^* \propto \partial \log(r)_d &\equiv \left. \frac{\partial \log(r)}{\partial d} \right|_{p_s=p_s^*, p_c=p_c^*, d=d^*}. \end{aligned}$$

305 By analyzing these three gradients, it is possible to determine the equilibria
 306 and stability conditions of the model.

307 **3.1. Equilibria and stability.** While this model has no true equilibria, be-
 308 cause it is stochastic, it does have steady state expected values for the state
 309 variables. It turns out that there are only two possible steady states in this

310 model. Let $\hat{d}, \hat{p}_s, \hat{p}_c$ denote expected values of the state variables that satisfy
 311 $\Delta p_s^* = \Delta p_c^* = \Delta d^* = 0$. Either the system comes to rest at a detection-
 312 less steady state where $\hat{d} = 0$ and $\hat{p}_s = u(s - k)/(k(1 - u))$, or it comes
 313 to rest where $\hat{d} > 0$ and $\hat{p}_s = 0, \hat{p}_c = 1$. In the Supporting Information, I
 314 show how to derive conditions for the stability of both the $\hat{d} = 0$ steady state
 315 and the $\hat{d} > 0$ steady state. In the remainder of this section, I present these
 316 conditions and try to motivate their logic.

317 3.1.1. *Condition for detection $d > 0$ to invade.* The condition for detection
 318 to invade from zero is:

$$\ell w_0 < p_s^*(k - su)(1 - f'(0)), \quad (7)$$

319 where $p_s^* = u(s - k)/(k(1 - u))$ is the probability of individual learning when
 320 $d^* = 0$, and $f'(0)$ is $\partial f / \partial d|_{d=0}$, the initial rate of increase in false alarms of
 321 environmental change. This expression confirms the intuition that detection
 322 can more easily invade when its direct fitness cost, ℓ , is low. Also intuitively,
 323 when false positives increase slowly with detection, $f'(0)$ is small, detection
 324 more easily invades.

325 3.1.2. *Level of detection when $\hat{d} > 0$.* When detection does invade and in-
 326 crease from zero, the learning state variables evolve to $\hat{p}_s = 0$ and $\hat{p}_c = 1$.
 327 There is no similarly simple expression for the value of \hat{d} . The expression for
 328 the steady-state amount of detection is complex, but is defined implicitly by:

$$f'(\hat{d}) = \frac{u(p_c^* - p_s^*)(s - k) - \ell w_0}{(1 - u)(p_c^* - p_s^*)(Q - (s - k))}, \quad (8)$$

329 or equivalently, using the fact that $f'(d) = \frac{\Delta f}{\Delta d}$:

$$\underbrace{\Delta d \left(u(p_c^* - p_s^*)(s - k) - \ell w_0 \right)}_{\text{marginal benefit}} = \underbrace{\Delta f (1 - u)(p_c^* - p_s^*)(Q - (s - k))}_{\text{marginal cost}}, \quad (9)$$

330 where Q has the same form as in Equation 6, quantifying exactly the expected
 331 probability of acquiring adaptive behavior via social learning, conditioned on
 332 the environment being stable ($t > 0$). Equation 9 states what in hindsight is
 333 obvious: selection converges to the value of d^* at which the marginal benefits
 334 of detection are equal to the marginal costs of false positives. But it also
 335 identifies the precisely relevant marginal benefits and costs, which I believe
 336 is less obvious. I'll unpack this equation one part at a time.

337 First, notice that everything except the direct cost of detection, ℓw_0 , is
 338 scaled by the term $p_c^* - p_s^*$. This difference is how much more individual
 339 learning is expressed when an individual believes the environment has just

340 changed ($t = 0$). When this difference is zero, detection has no effect on
 341 behavior, because learning is not contingent upon the signal. The direct cost
 342 ℓw_0 is the unconditional marginal cost of investing in detection. It is unaf-
 343 fected by the difference $p_c^* - p_s^*$, because it is a flat fitness cost that always
 344 reduces net benefits.

345 The left side of Equation 9 summarizes the net marginal benefits of detec-
 346 tion. The probability that the environment changes (u) is multiplied by the
 347 expected net benefit of individual learning, $s - k$. This is the rate of ben-
 348 efit from correct detection. This expected benefit is unconditional on the
 349 frequencies of individual and social learning, because when $t = 0$, social
 350 learning never pays and individual learning's fitness is independent of the
 351 frequencies p_s^*, p_c^* . Finally, the marginal cost of detection ℓw_0 is subtracted
 352 to yield the net benefit of detection.

353 The right side of Equation 9 is the marginal cost of detection. This expres-
 354 sion quantifies the expected foregone benefits of mistakenly learning individ-
 355 ually when the environment is stable (a false positive). Q is the probability of
 356 acquiring adaptive behavior via social learning (given that $t > 0$), and $s - k$ is
 357 again the net benefit of individual learning. The difference is the net benefit
 358 of social learning, or rather here the net cost of a false positive, which induces
 359 an individual to learn individually when it might have learned socially.

360 3.1.3. *Condition for detection $d^* > 0$ to be stable.* Whether or not Express-
 361 sion 7 is satisfied, it is possible for detection to be stable once common. The
 362 condition in this case is a complicated expression that yields little qualitative
 363 insight, but I show in the Supporting Information that it can be satisfied even
 364 when detection cannot invade.

365 To begin to understand why this is the case, note that there are three dy-
 366 namic regimes in the model. Figure 2 illustrates these. Each plot shows the
 367 phase plane dynamics of p_s^* and d^* when $p_c^* = 1$. The state variable p_c^* can
 368 be fixed at one, because it evolves to one very quickly for most parameter
 369 combinations. This allows us to understand the reduced two-dimensional
 370 system, as shown in the figure. In each plot, the gray lines with arrows show
 371 the flow of the system at each point in the p_s^*, d^* space. The black curve is the
 372 p_s^* fitness isocline, the combinations of p_s^*, d^* that satisfy $\partial \log(r)_{p_s} = 0$. The
 373 red curve is the d^* fitness isocline. Above the black curve, selection decrease
 374 p_s^* . Below the black curve, selection increases p_s^* . Above the red curve, se-
 375 lection decreases d^* , and below it selection increases d^* . The false positive
 376 function is set to $f(d) = ad/(1 + a - d)$.

377 Now consider each plot in Figure 2 in turn. First, when ℓ is very small,
 378 in panel (a), detection can both invade from zero and is stable once large.
 379 Detection invades at the point where the black curve meets the bottom axis.
 380 Since this is below the red curve in (a), selection increases detection. In this




FIGURE 2. Dynamics of signal detection, as a function of the cost of detection, ℓ . Each figure plots the dynamics in the trait space p_s^* and d^* . $p_c^* = 1$ in each case, which clarifies the presentation, without loss of generality. The gray streams represent the evolutionary flow of these characters. The black curve is the combinations of p_s^* , d^* at which there is no directional change in p_s^* . The red curve is the combinations of p_s^* , d^* at which d^* does not change. The three panels vary in the direct cost of detection, ℓ , while holding constant $b = 0.1$, $u = 0.3$, $k = 0.35$, $w_0 = 1$, $s = 1$, $a = 0.1$. (a) $\ell = 0.02$: There is only one equilibrium here, where the red curve meets the left margin, at $p_s^* = 0$. (b) $\ell = 0.08$: There is now an unstable internal equilibrium, where the red and black curves intersect, and two stable points, at the left end of the red curve and the bottom end of the black curve. (c) $\ell = 0.20$: The only equilibrium in this case is where the black curve meets the bottom margin, where $p_s^* = u(s - k)/(k(1 - u)) \approx 0.8$ and $d^* = 0$.

381 case, detection will always evolve to $\hat{d} > 0$ and $\hat{p}_s = 0$, $\hat{p}_c = 1$. Second, when
 382 ℓ is intermediate, as in panel (b), detection cannot invade from zero but can
 383 be stable once large. In this case, detection may come to rest at $\hat{d} = 0$ or
 384 $\hat{d} > 0$, depending upon initial conditions. Third, if ℓ is sufficiently large, as
 385 in panel (c), detection can neither invade nor be stable. In this case, detection
 386 will always remain at $\hat{d} = 0$.

387 Another way to summarize the same dynamic is to plot the best response
 388 values of p_s , p_c and d as a function of d^* . A best response here is the value
 389 of the trait that will maximize fitness, conditioned on the value of the other




FIGURE 3. Best response values of p_s, p_c, d , as a function of d^* and the cost of detection, ℓ . In each plot above, the green, red, and black curves are the respective values of p_s, p_c, d that jointly maximize fitness, given the value of d^* on the horizontal axis. Where the black curve is above the diagonal, larger values of d increase fitness. Below the diagonal, smaller values of d increase fitness. Where the black curve crosses the diagonal is an equilibrium. Parameter values held at $b = 0.1, u = 0.2, k = 0.3, s = 1, w_0 = 1, a = 0.5$. Plots (a), (b) and (c) vary ℓ so as to illustrate the same three regimes as in Figure 2. (a) Low cost and globally stable detection where the black curve crosses the diagonal. (b) Intermediate cost and a bi-stable regime. (c) High cost and globally stable $d^* = 0$.

390 traits. Figure 3 shows these best responses. In each plot, the green, red and
 391 black curves plot the respective values of p_s, p_c, d that maximize fitness, given
 392 a population with common trait value d^* on the horizontal axis. I compute
 393 these by allowing p_s and p_c to go to their equilibrium values, given d^* . This
 394 provides the values for the green (p_s) and red (p_c) curves. Then I compute
 395 the fitness maximizing value of d , conditioned on d^* and the best response
 396 values of p_s and p_c . So to see what values of p_s, p_c, d are favored when d^* takes
 397 a particular value, find the value of d^* on the horizontal axis and then go up
 398 to find the values of p_s (green curve), p_c (red curve), and d (black curve) at
 399 that point.

400 These plots clearly illustrate a key result of the model. The stable equilib-
 401 rium for $d^* = \hat{d}$, found where the black curve crosses the diagonal, always
 402 occurs where $p_s = 0, p_c = 1$. Behavior is perfectly separated by receiving the
 403 signal. Why should this be the case? Why can't d^* stabilize where either p_s




FIGURE 4. Mean fitness at the detection steady state. In each plot, the dashed gray, light red, and light green curves are the best-response d , p_c , and p_s , corresponding to the curves in Figure 3. The solid black curve is the mean fitness at d^* on the horizontal axis. The solid blue curve is the quality of social information, Q . The solid orange curve is the average rate of individual learning. The points on each curve show the values at the $\hat{d} > 0$ detection steady state. (a) Low cost of detection and mean fitness at $\hat{d} > 0$ is greater than at $\hat{d} = 0$. (b) Higher detection cost. Mean fitness is still higher at $\hat{d} > 0$, but it declines initially during invasion. (c) At very high detection cost, mean fitness at $\hat{d} > 0$ can be lower than at $\hat{d} = 0$. In all three plots, $u = 0.3$, $k = 0.4$, and $a = 0.1$. Dynamics are explained in the main text.

404 or p_c is intermediate between zero and one? The reason is complex enough
 405 to deserve it's own section, to follow.

406 **3.2. Mean fitness and the dynamics of detection.** The mean fitness (log-
 407 geometric growth rate of the common type) in the population either remains
 408 constant or decreases during invasion. But near the $\hat{d} > 0$ steady state, mean
 409 fitness may be both greater than or less than fitness at $\hat{d} = 0$. As a result, once
 410 detection evolves, the population could be either better or worse off than if
 411 no one bothered to detect environmental change. In this section, I attempt
 412 to explain these dynamics. In the process, it will become clear why $d^* > 0$
 413 always stabilizes where $p_c^* = 1$ and $p_s^* = 0$.

414 Figure 4 illustrates the dynamics of mean fitness, as detection invades and
 415 stabilizes. Each plot in this figure has the same axes as the plots in Figure 3.
 416 In all three plots, $u = 0.3$, $k = 0.4$, and $a = 0.1$, chosen for clarity of

417 presentation. The dashed curves represent the same kind of best response
 418 profiles as in Figure 3, with gray for d , light red for p_c , and light green for p_s .
 419 The solid curves now represent mean fitness (black), the rate of individual
 420 learning (orange), and the quality of social information (blue). The orange
 421 curves are rates of individual learning, $uE_c^* + (1 - u)E_s^*$, where $E_c^* = d^*p_c^* +$
 422 $(1 - d^*)p_s^*$ is the probability of individual learning when the environment
 423 has just changed ($t = 0$) and $E_s^* = f^*p_c^* + (1 - f^*)p_s^*$ is the probability
 424 when the environment is stable ($t > 0$). The blue curves are Q as defined
 425 earlier. The black mean fitness curve is just the functional component of
 426 fitness, $I + S - \ell d^*$, as in Equation 4. Each of these solid curves is computed
 427 for a population with d^* on the horizontal axis and p_c^* and p_s^* from the best
 428 response curves.

429 In the Supporting Information, I present mathematical analysis of rates
 430 of change of individual learning (orange) and quality of social information
 431 (blue) as they respond to changes in detection d and individual learning p_s .
 432 Here, I provide a verbal summary to motivate understanding of how the av-
 433 erage rate of individual learning and Q contribute to the dynamics of mean
 434 fitness.

435 When detection increases from zero, the mutant individual's rate of indi-
 436 vidual learning increases. This increase results from more adaptive individ-
 437 ual learning at $t = 0$. But it also results from more maladaptive individual
 438 learning at $t > 0$, because of false positives f . Selection then favors a reduc-
 439 tion in p_s , to both compensate for the false positives as well as reap greater Q
 440 resulting from the spread of detection in the population. But this reduction
 441 in p_s , once it spreads through the population, reduces both the overall rate
 442 of individual learning (orange) as well as the quality of social information
 443 Q (blue) in the population. This reduction in the quality of social informa-
 444 tion cancels any mean fitness benefit of detection. So during this phase of
 445 the dynamics, the rate of individual learning in the population declines, but
 446 accompanying decline in the quality of social information means the popula-
 447 tion experiences no average fitness benefit from avoiding individual learning.
 448 This dynamic is very much like the one that generates constant mean fitness
 449 in many models of this kind (Boyd and Richerson 1995).

450 However, mean fitness is not always constant during this phase of the dy-
 451 namics. If detection is individually costly ($\ell > 0$), as in the middle plot in
 452 Figure 4, mean fitness (black curve) will actually decline as detection invades.
 453 This decline does not prevent detection from invading, however, because the
 454 mutant is playing the market and does experience a relative fitness benefit
 455 initially. It is only once the rest of the population catches up that the qual-
 456 ity of social information declines and reduces mean fitness. Thus in such
 457 scenarios, detection invades but actually makes the population worse off.

458 Once detection is high enough, however, p_s^* reaches zero and cannot be
 459 reduced any further. Any further increases in detection now increase the
 460 rate of individual learning, because false positives at $t > 0$ cannot be com-
 461 pensated for by reducing p_s . The orange curves rise. As a result, the quality
 462 of social information also rises, producing a population benefit of more ac-
 463 curate behavior to imitate. Detection d^* can continue to increase, raising
 464 both the quality of social information and mean fitness, until the rate of false
 465 positives (f^*) satisfies Equation 8. At that point, selection stabilizes $\hat{d} > 0$.

466 Mean fitness at $\hat{d} > 0$ can therefore be higher than that at $\hat{d} = 0$, as in the
 467 lefthand and middle plots in Figure 4. But if the costs of detection ℓ are large
 468 enough, as in the righthand plot, then mean fitness at $\hat{d} > 0$ may actually be
 469 lower than that at $\hat{d} = 0$. The same dynamic as above is at work here, but now
 470 the direct cost of detection is large enough that the increase in mean fitness
 471 near $\hat{d} > 0$ cannot overcome it.

472 This dynamic helps to understand why $\hat{d} > 0$ can only stabilize where
 473 $\hat{p}_c = 1$ and $\hat{p}_s = 0$ and why mean fitness can only increase in the same
 474 region. Until p_s^* reaches its minimum, selection can compensate for an in-
 475 crease in false positives by both reducing individual learning when $t > 0$.
 476 This compensation has the consequence of spoiling the quality of social in-
 477 formation, erasing fitness gains of invaders. This is similar to the dynamic in
 478 Rogers' model and many similar models (Rogers 1988, Boyd and Richerson
 479 1995), in which invading social learners eventually spoil the quality of social
 480 information, erasing any fitness gains for the population. However in this
 481 model, once p_s^* cannot be further reduced to compensate for increasing f ,
 482 invaders improve the quality of social information. Now mean fitness (black
 483 curves in Figure 4) will rise, both because of (1) the direct benefits of detec-
 484 tion allowing individuals to allocate expensive individual learning to when
 485 it is needed most and (2) the population side effect of improving the quality
 486 of social information.

487 Further increase in d^* beyond $\hat{d} > 0$ would increase mean fitness, as can
 488 be seen by the location of the black points in Figure 4 always lying to the left
 489 of the peak of the black curve. However, in the absence of some factor like
 490 kin selection (in a non-viscous population structure), natural selection will
 491 not maximize mean fitness in this model. Detection is individually costly, but
 492 produces a population benefit by increasing the quality of social information,
 493 near steady state. This is a kind of collective action dilemma, similar to the
 494 basic collective action dilemma embodied in individual and social learning:
 495 individual learning is individually costly, but produces population benefits.
 496 Ironically, kin selection would increase detection at $\hat{d} > 0$, but also narrow
 497 the conditions that allow detection to invade, because now the depression of

498 mean fitness during invasion ($\ell > 0$, as in the middle plot) would reduce
499 inclusive fitness.

500

4. DISCUSSION

501 **4.1. When does detection evolve?** Detection of recent environmental change
502 evolves under many circumstances that favor reliance on social learning at
503 all. Selection further favors detection when individual learning is costly rela-
504 tive to the rate of environmental change, detection is efficient (not too many
505 false positives), and detection is not too costly. These conclusions can be
506 read directly from Condition 7. To understand them, consider that the initial
507 problem detection solves is allocation of costly individual learning to gener-
508 ations immediately following environmental change. Detecting change al-
509 lows individuals to likewise allocate more social learning to generations in
510 which the environment is stable. But if individual learning is cheap relative
511 to the rate of environmental turnover, then little social learning is favored
512 even when the environment is stable. At the extreme where the rate of change
513 u is larger than the effective costs of individual learning k/s , detection can
514 never invade the population.

515 However, detection does invade over a broad range of parameter values.
516 This is most evident perhaps in the sensitivity plots in the Supporting Infor-
517 mation (Figures 5 and 6). The intuition behind this result is that, whenever
518 substantial social learning is favored in the absence of detection, as it rou-
519 tinely is in such models, there will be a basic allocation problem that can
520 be addressed by detecting recent environmental change. Social learning is a
521 risky, high variance learning strategy, relative to individual learning. Just af-
522 ter a change in the environment, all social learning results in zero probability
523 of acquiring adaptive behavior. This effect is very stark in this kind of model,
524 because all adaptive information is lost when the environment changes. But
525 the general principle appears robust, as it remains even in cumulative culture
526 models in which some adaptive behavior can persist (McElreath 2010, e.g.).
527 Detection reduces the fitness variance of social learning, by allocating more
528 of it to when it is safest to use.

529 **4.2. How does detection work?** Of course these results must overstate the
530 probability of detection's evolution, because the model assumes a cue of en-
531 vironmental change is available and that the organism can discover it. The
532 nature of such cues is left abstract in the model, but the sensory abilities of the
533 organism and structure of the population must constrain the possibilities.
534 Potentially general cues of recent environmental change may include poor
535 health or fertility of conspecifics. Organisms with sophisticated communi-
536 cation, like humans, may also detect change by paying attention to reflections

537 of older individuals. For example, contemporary arctic peoples possess elab-
538 orate, ecologically-accurate models of their environments. When events fall
539 outside past patterns, younger individuals can and do benefit from listening
540 to such observations (Fox 2002, Weatherhead *et al.* 2010). Under this view
541 of detection, the costs are in attention and processing, being possibly quite
542 small.

543 Another idea is that naive individuals learn socially first and that the cost of
544 detection in this model represents the cost of trying out socially-learned be-
545 havior. This would make detection here like “critical social learning” (Boyd
546 and Richerson 1996, Enquist *et al.* 2007), in which individuals first learn so-
547 cially, have a chance of diagnosing maladapted behavior, and can finally use
548 individual learning as a last resort. Like detection in this model, critical so-
549 cial learning can be stable even when it cannot invade, and it can raise mean
550 fitness. Unlike detection however, critical social learning models currently
551 contain no false positives—although there is transmission error, critical so-
552 cial learners never mistake adaptive behavior for maladaptive. The detection
553 model here in contrast contains inherent risk of excessive individual learn-
554 ing, because of mistaken diagnosis. But the key point is that critical social
555 learning may be an alternative mechanism for detecting temporal environ-
556 mental change, one that is rather accessible to evolution. Its dynamics will be
557 different from the use of cues specifically tied to ecological change, such as
558 those mentioned in the previous paragraph, but may nevertheless coevolve.

559 **4.3. What are detection’s effects?** Once detection does invade, the distri-
560 bution of individual and social learning through time changes. A dominant
561 result is that a population that invests in detecting recent change will ex-
562 hibit less individual learning overall. It will also exhibit much less individual
563 learning during periods of environmental stability.

564 This allocation of individual learning to periods just following a change
565 results in a rapid increase in the frequency of adaptive behavior, just after a
566 change in environment. But it also results in a very slow increase afterwards.
567 As a result, the frequency of adaptive behavior in stable environments may
568 not look very different, after detection evolves. However, the frequency of
569 adaptive behavior recently following a change in the environment will look
570 quite different, showing a rapid increase. In the end, a snapshot of a popu-
571 lation in which detection has evolved will show a higher reliance on social
572 learning compared to what evolves in the absence of detection.

573 All of these population dynamics combine to allow average fitness, or the
574 expected population growth rate, to rise after detection evolves. This increase
575 in mean fitness is usually quite modest in this model, much smaller than
576 that demonstrated from cumulative culture models. However the increase
577 appears for much the same reason: detection allows individual and social

578 learning to work together synergistically, rather than competitively (Boyd
579 and Richerson 1995). Learning individually just after a change in the envi-
580 ronment both helps those doing the learning as well as the social learners in
581 following time periods. However, the benefit to subsequent generations does
582 not lead to natural selection reducing individual learning, because the ben-
583 efit was produced at a time that social learning produces no fitness increase
584 at all. Therefore the “freeloading” by social learners during stable periods
585 does not threaten to erode the public good provided by individual learners
586 just after a change. But if detection is very costly or individual learning is
587 sufficiently cheap, then the change in mean fitness may be very small or even
588 negative.

589 **4.4. Model variations.** I have chosen this model’s features because they rep-
590 resent a central case for analysis, one comparable to existing theory. How-
591 ever, all models are necessarily special, and so future work should address
592 other model assumptions.

593 *Spatial environmental variation.* It has long been recognized (Hedrick *et al.*
594 1976) that spatial and temporal environmental variation produce different
595 selection regimes. This is just as true for gene-culture coevolutionary mod-
596 els (McElreath *et al.* in press, Nakahashi *et al.* in press). “Detection” in a
597 spatial variation context would mean the population is sub-divided into a
598 number of patches. A different behavior is optimal in each patch. Individ-
599 uals can evolve different learning strategies depending upon whether or not
600 they are recent immigrants to a local patch. A first conjecture is that selection
601 will favor greater reliance on social learning for recent immigrants. It would
602 also be possible to examine whether selection may favor residents’ ignoring
603 immigrants, when choosing models to learn from. An important question
604 to ask of such a model is whether adjusting use of social learning depend-
605 ing upon migration status will allow unbiased social learning to maintain
606 cultural variation, even when it cannot in traditional models.

607 Spatial and temporal variation may also interact in unanticipated ways.
608 Such interactions have been well-explored in the study of dispersal (Schreiber
609 2010, for a recent example), but the importance of these phenomena is po-
610 tentially much more general (Williams and Hastings 2011).

611 *Other learning strategies.* Unbiased social learning, in which a single tar-
612 get of learning is chosen independent of its behavior, is a very special case.
613 The most-discussed alternatives includes conformist transmission (Boyd and
614 Richerson 1985, Henrich and Boyd 1998) and some kind of payoff or success
615 or prestige biased transmission (Boyd and Richerson 1985, Henrich 2001,
616 McElreath *et al.* 2008).

617 In the case of conformist transmission, recent debates over whether or
618 not selection will favor it provide a natural opening to consider contingent
619 use. Conformist transmission was originally studied as an adaptation to spa-
620 tial environmental variation (Boyd and Richerson 1985, Henrich and Boyd
621 1998). Wakano and Aoki (2007) later studied a model of conformist learning
622 in which there was only temporal variation, finding that the conditions that
623 favored conformist transmission were very restrictive. Some of the same au-
624 thors have more recently studied conformist transmission under both tem-
625 poral and spatial variation, confirming the original intuition that it is an
626 adaptation to spatial variation (Nakahashi *et al.* in press). McElreath *et al.*
627 (in press) have recently shown that a mix of temporal and spatial variation
628 can also favor a strong reliance on conformist transmission, as in Henrich
629 and Boyd's model. Finally, a recent Bayesian model (Perreault *et al.* in press)
630 demonstrates a robust reliance on conformist transmission, even when the
631 environment varies only temporally.

632 An explicit consideration of contingent use of conformist transmission as
633 a function of cues of environmental change and migration status should help
634 to unify this literature. It would also help in interpretation of experimental
635 results. All of the existing experimental and quasi-experimental studies of
636 social learning contained analogues of only temporal environmental varia-
637 tion. While conformist transmission has been found in some of these cases
638 (Kameda and Nakanishi 2002, McElreath *et al.* 2005, 2008), it has not always
639 been found (Eriksson *et al.* 2007, Eriksson and Coultas 2009). Experiments
640 that allow for the analogue of spatial variation should provide cleaner tests.

641 *Learning costs.* When there are multiple domains of behavior to be learned,
642 and the costs of learning vary among them, how will selection design learn-
643 ing? Since the problems the organism needs to solve may change across space
644 and time, it is problematic to assume that there can be a genetic locus control-
645 ling reliance on social learning in each domain. Should an organism attempt
646 to estimate a cost of individual learning in each domain, or rather adapt to a
647 fitness-weighted average of the domains?

648 5. CONCLUSION

649 To return to the problem that motivated this model: How do these re-
650 sults reflect on the interpretation of social learning experiments? If human
651 or other animal participants do have strategies attuned to cues of environ-
652 mental change, we will need to consider whether or not our experiments ac-
653 cidentally include too many or too few such cues. For example, in the typical
654 experiment, all participants are equally naive at the start. This may function

655 as a social cue of recent environmental change, in that it favors increased re-
656 liance on individual learning. On the other hand, experiments that provide
657 participants with no way to detect changes in the underlying payoffs may
658 accidentally provide cues of environmental stability. In the end, explicitly
659 designing both laboratory and field studies with contingent strategy use in
660 mind will provide clearer tests of theory.

SUPPORTING INFORMATION

661

662 **Derivation of q_t .** First, note that just after a change in the environment, q_t
 663 resets to $q_0 = 0$. One generation after a change in the environment, $t = 1$,
 664 the expected chance of acquiring adaptive behavior via social learning is:

$$q_1 = d^* p_c^* s + (1 - d^*) p_s^* s.$$

665 Here's how to motivate the expression above. The only way to acquire adap-
 666 tive behavior this soon after a change in the environment is to target an in-
 667 dividual who learned individually in the last generation. A proportion $d^* p_c^*$
 668 of the previous generation correctly noticed the change in the environment
 669 and chose to learn individually with chance p_c^* . The remaining proportion
 670 $1 - d^*$ failed to detect the change in the environment and continued to learn
 671 individually with chance p_s^* .

672 Next, in any generation $t > 1$, the quality of social information is given by:

$$q_t = (1 - f^*) (p_s^* s + (1 - p_s^*) q_{t-1}) + f^* (p_c^* s + (1 - p_c^*) q_{t-1}). \quad (10)$$

673 Here's how to motivate this expression. A proportion $1 - f^*$ of the population
 674 did not wrongly conclude that the environment changed recently. They learn
 675 individually p_s^* of the time and update socially $1 - p_s^*$ of the time. Social
 676 updating leads to successful acquisition of adaptive behavior with probability
 677 q_{t-1} . The rest of the population, a proportion f^* , thinks the environment just
 678 changed and updates accordingly.

679 The above recurrence equation for q_t (Equation 10) can be solved explicitly
 680 for a function q_t that is not a function of q_{t-1} . It is a linear recurrence and so
 681 several methods exist. I used the *Ansatz* method of guessing the form and
 682 proving it was correct. The resulting function is:

$$q_t = s \left(1 - (1 - d^* p_c^* - (1 - d^*) p_s^*) (1 - f^* p_c^* - (1 - f^*) p_s^*)^{t-1} \right),$$

683 for all $t \geq 1$.

684 **Invasion and stability conditions for any function $f(d)$.** To find the con-
 685 ditions for this equilibrium to exist and be stable, we can observe that the
 686 dynamics of d^* , p_s^* are governed by two null clines, where the change in each
 687 state variable is zero, as a function of p_s^* (Figure 2). Both null clines are con-
 688 vergent, in the sense that dynamics take each state variable closer to its null
 689 cline. Thus if d^* is plotted on the vertical axis and p_s^* on the horizontal axis,
 690 then the system moves up when it is below the null cline for d^* and left when
 691 it is to the right of the null cline for p_s^* . Because these null clines cross only
 692 once, we can inspect the four end points on both the left and bottom axes to
 693 summarize the dynamics of the system.

694 Consider first the bottom axis, where $d^* = 0$. The p_s^* null cline will always
 695 intersect the bottom axis at $p_s^* = u(s - k)/(k(1 - u))$. Then depending upon
 696 whether the d^* null cline lies left or right of this point determines whether
 697 detection can invade. The null cline for d^* must lie to the right of the null
 698 cline for p_s^* , for d^* to increase. Otherwise the system at the invasion point
 699 $p_s^* = u(s - k)/(k(1 - u))$ will be above the d^* null cline and decrease. If
 700 instead the d^* null cline is right of $p_s^* = u(s - k)/(k(1 - u))$, the system will
 701 lie below the d^* null cline and therefore d^* will increase. So the condition
 702 for d^* to increase from the point $d^* = 0, p_s^* = u(s - k)/(k(1 - u)), p_c^* = 1$
 703 is given by asking when the value of p_s^* that makes $\partial \log(r)_d|_{d^*=0} = 0$ is
 704 greater than $p_s^* = u(s - k)/(k(1 - u))$. When instead this point is lower
 705 than $p_s^* = u(s - k)/(k(1 - u))$, detection cannot invade from zero. Reducing
 706 this condition tells us that stability at $d^* = 0$ requires either that $k \leq su$ or
 707 $s \leq k$, or when $k > su$ and $s > k$, it requires:

$$\ell > \frac{(s - k)u(k - su)(1 - f'(0))}{k(1 - u)w_0}. \quad (11)$$

708 When this condition is satisfied, detection cannot increase from zero.

709 The second condition for the internal unstable equilibrium to exist is that
 710 the null cline for d^* , along the left axis where $p_s^* = 0$, be greater than the null
 711 cline for p_s^* (Figure 2). This reduces to the condition:

$$\ell < \frac{(s - k)u}{w_0} \quad (12)$$

712 and

$$\ell < \frac{(s - k)u}{w_0} - f'(d_1) \frac{k(1-u)((s-k)(f_2(1-u)+u)^2 - (f_1-f_2)(1-u)(k(f_2(1-u)+u)-s))}{(f_1(1-u)+u)(k(f_2(1-u)+u)-s)w_0}, \quad (13)$$

713 where d_1 is the value of d^* that satisfies $\partial \log(r)_d|_{p_s^*=0, p_c^*=1} = 0$ and d_2 the
 714 value of d^* that satisfies $\partial \log(r)_{p_s}|_{p_s^*=0, p_c^*=1} = 0$. The symbol $f_1 \equiv f(d_1)$ and
 715 $f_2 \equiv f(d_2)$. For condition 13 to be less than condition 12, it is also necessary
 716 that:

$$s < k(f_2(1 - u) + u).$$

717 In summary, when detection cannot invade (condition 11 is satisfied) but
 718 is stable when large (condition 13 is satisfied), the dynamics contain an inter-
 719 nal unstable equilibrium. This proves that the signal detection equilibrium
 720 with stable $d^* > 0$ and $p_s^* = 0, p_c^* = 1$ is stable for a broader range of values
 721 than will allow $d^* > 0$ to invade from $d^* = 0$.

722 The dynamics of the system can also be summarized in terms of these two
 723 conditions. There are three possible combinations. First, detection can in-
 724 vade when rare and be stable when common. This requires that condition 11
 725 be false and condition 13 true. Second, detection cannot invade when rare
 726 but can be stable once at a large enough value. This holds when condition 11
 727 is true and condition 13 is true. Third, detection can neither invade nor be
 728 stable once large. This holds when condition 11 is true and condition 13 is
 729 false.

730 **Sensitivity plots.** It is much easier to appreciate the effects of the parame-
 731 ters on invasion, stability, and equilibrium detection rate by using sensitiv-
 732 ity plots. In Figures 5 and 6, I illustrate how the parameters influence the
 733 relative sizes of the invasion and stability conditions, as well as the steady
 734 state detection rate. These plots all use the box hyperbola ROC function,
 735 $f(d) = ad/(1 + a - d)$.

736 Figure 5 plots evolutionary outcomes of the model, within the parameter
 737 space defined by k , the cost of individual learning, and u , the instability of
 738 the environment. Each of the six plots varies the accuracy of detection, a ,
 739 and the direct cost of detection, ℓ . All other parameters are held constant at
 740 $w_0 = 1, s = 1$. The red regions enclose all combinations of k, u that lead
 741 d^* to increase from zero. These are the invasion regions. The shaded re-
 742 gions enclose all combinations of k, u for which detection can be stable once
 743 large enough. These are the stability regions. The degree of shading in the
 744 stability regions represents the amount of detection at equilibrium, for each
 745 combination of k, u . Pure black represents $\hat{d} = 1$ while pure white represents
 746 $\hat{d} = 0$.

747 In every case, the red invasion region does not extend above the diagonal
 748 where $k/s = u$. When $u > k/s$ and $d^* = 0$, individuals are already using
 749 individual learning 100% of the time. If individuals had a flawless signal of
 750 environmental change, then selection would favor detection and using so-
 751 cial learning when the environment is stable. But the signal is never perfect.
 752 Instead, attempts to detect stability always lead to some erroneous decisions
 753 to learn socially. When u is large, the probability the environment has not
 754 changed will be small and comparable to the rate of mistakes in concluding
 755 that the environment has not changed. Since individual learning is so cheap,
 756 when $u > k$, the risks do not outweigh the costs, and detection can never
 757 invade.

758 For combinations of small u and large k (lower-right corner of each plot),
 759 invasion is similarly impossible. In these regions, very little individual learn-
 760 ing is favored, because of its high cost and the infrequency of change in the




FIGURE 5. Sensitivity of invasion and stability to the parameters k , the cost of individual learning, and u , the probability of environmental change. In each plot, the red boundary contains the combinations of k, u that allow detection to invade. The shaded region shows all combinations of k, u at which detection is stable when common. Darker shading indicates higher equilibrium detection, \hat{d} , with pure black representing $\hat{d} = 1$ and pure white $\hat{d} = 0$. Inside the shaded region, $\hat{p}_s = 0, \hat{p}_c = 1$. Outside the shaded region, $\hat{p}_s = u(s - k)/(k(1 - u)), \hat{d} = 0$. Note that detection can be stable at high values even when there are no combinations of k, u that allow invasion, as in panel (c).

761 environment. For example, where $u = 0.1$ and $k = 0.9$, the stable proportion of individual learning is $p_s^* \approx 0.01$. Now the environment doesn't
 762 proportion of individual learning is $p_s^* \approx 0.01$. Now the environment doesn't
 763 change much, and attempts to detect change will generate errors at a rate
 764 comparable to the probability of true change.

765 Detection can be stable over a larger region, however, above $u = k/s$.
 766 Even though detection cannot invade where $u > k/s$, it may be stable on
 767 large. Once detection is accurate enough, it allows individuals to allocate
 768 more individual learning to when it is most needed and simultaneously re-
 769 duce their overall reliance on individual learning, which is expensive.

770 A similar phenomenon does not appear in the corner where k is large and
 771 u is small, because so little individual learning is favored there that detection
 772 cannot reduce fitness costs much by reducing overall reliance on individual
 773 learning—there just isn't much reliance to reduce. As a result, the stability
 774 regions can extend far above the $k/s = u$ invasion boundary, but not very
 775 far beyond the lower-right of it.

776 In the top row of Figure 5, accuracy of detection is set fairly low, to $a =$
 777 0.5 . At this accuracy, in order to correctly detect 50% of all environmental
 778 changes, an individual would suffer false alarms 25% of the time. In order to
 779 correctly detect 90% of changes, an individual would suffer false alarms 75%
 780 of the time! The three plots in this row vary the direct cost of detection, ℓ .
 781 Both the invasion and stability regions shrink rapidly with increases in the
 782 direct cost of detection.

783 In the bottom row of Figure 5, accuracy of detection is set fairly high, to
 784 $a = 0.1$. And at this accuracy, 50% true detection implies a false alarm rate
 785 of only 8.3%. A 90% detection rate implies a 45% false alarm rate. When
 786 accuracy is this high, changes in the direct cost of detection, ℓ , have much
 787 less of an effect on the stability region.

788 It is easier to appreciate the effect of ℓ by holding u constant and varying ℓ
 789 and k . I do this in Figure 6, for three values of u (0.05, 0.20, 0.40) and the same
 790 two values of a as in Figure 5 (0.5, 0.1). In the space defined by ℓ and k , the
 791 red invasion region rises above zero where $k = u$. This is the boundary on
 792 the diagonal in Figure 5. We can now see, however, that the invasion region
 793 extends all the way to the right, provided that $\ell = 0$. For $\ell > 0$, intermediate
 794 values of k have the largest invasion potential. This reflects the same tradeoffs
 795 that I explain for the u, k parameter space.

796 The effect of the accuracy of individual learning, s , is to compress the space
 797 defined by k . The true dimension of the cost of individual learning is k/s , not
 798 k . When $s = 1$, as in the previous figures, k summarizes the cost of individ-
 799 ual learning. But for smaller s , the horizontal axis is effectively compressed,
 800 otherwise leaving the geometry unchanged.

801 **Dynamics of L and Q .** The expected rate of individual learning, as a func-
 802 tion of d, p_c, p_s is:

$$L = u(dp_c + (1 - d)p_s) + (1 - u)(fp_c + (1 - f)p_s).$$




FIGURE 6. Sensitivity of invasion and stability to the parameters k , the cost of individual learning, and ℓ , the direct cost of detecting environmental change. In each plot, the red boundary contains the combinations of k, ℓ that allow detection to invade. The shaded region shows all combinations of k, ℓ at which detection is stable when common. Darker shading indicates higher equilibrium detection, \hat{d} , with pure black representing $\hat{d} = 1$ and pure white $\hat{d} = 0$. Note the scale of the vertical axis, which unlike the horizontal, only extends to 0.20.

803 This expression is for a mutant, but since individual learning is asocial, the
 804 rate for the population is analogous, using d^*, p_c^*, p_s^* .

805 We want to prove that L increases with d , in order to demonstrate that
 806 increasing detection increases individual learning on average. The rate of

807 change in L as a function of d is, via the chain rule:

$$\begin{aligned} \frac{dL}{dd} &= \frac{\partial L}{\partial d} + \frac{\partial L}{\partial f} \frac{df}{dd}, \\ &= (p_c - p_s) \left(u + (1 - u) \frac{df}{dd} \right). \end{aligned}$$

808 Since $df/dd > 0$ and $p_c - p_s > 0$, as conditions for $d > 0$ to invade in the
809 first place, $dL/dd > 0$.

810 Now consider the rate of change in L as a function of p_s . This is:

$$\frac{dL}{dp_s} = 1 - (1 - u)f - ud.$$

811 And this is also always positive, for any $0 < u < 1$ and $0 < f < d < 1$.
812 Increasing p_s increases L , and so decreasing p_s decreases L , explaining the
813 negative trend for L as p_s approaches zero.

814 Now consider the change in the quality of social information, Q . Again
815 via the chain rule, the rate of change in Q as a function of d^* is:

$$\begin{aligned} \frac{dQ}{dd^*} &= \frac{\partial Q}{\partial d^*} + \frac{\partial Q}{\partial f^*} \frac{df^*}{dd^*}, \\ &= su(p_c^* - p_s^*) \times \\ &\quad \frac{(1 - u)(p_s^* + f^*(p_c^* - p_s^*) + \frac{df^*}{dd^*}(1 - d^*(p_c^* - p_s^*) - p_s^*)) + u}{((1 - u)(f^*(p_c^* - p_s^*) + p_s^*) + u)^2}. \end{aligned}$$

816 While this appears complicated, it is always positive for any $0 < f^* < d^* <$
817 1 and $p_c^* > p_s^*$. It's worth noting also that the above is proportional to u ,
818 because detection improves Q partly by focusing more individual learning
819 on time periods where $t = 0$. This increases Q for all $t > 0$, as a consequence.
820 The more common change in the environment, the more detection helps Q .

821 Q also increases with p_s^* . The rate of change is:

$$\frac{dQ}{dp_s^*} = \frac{su(1 - (1 - u)f^*(1 - p_c^*) - d^*(u + (1 - u)p_c^*))}{(f^*(p_c^* - p_s^*)(1 - u) + p_s^*(1 - u) + u)^2}.$$

822 And this is also positive, for all $0 < f^* < d^* < 1$. Therefore as p_s^* decreases,
823 Q decreases.

824 In this model, the evolution of detection of environmental change can
825 both increase or decrease mean fitness. The analytical conditions for these
826 outcomes are complex functions of every variable in the model. However,
827 considering the special case where $\ell = 0$ does provide some qualitative in-
828 sight. The condition for expected fitness at the $\hat{d} > 0$ steady state to exceed

829 fitness at $\hat{d} = 0$ is:

$$\frac{k}{s} > \frac{u(1 - \hat{d})}{(1 - (1 - u)(1 - \hat{f})) (1 - (1 - u)\hat{f} - u\hat{d})}.$$

830 As the costs of individual learning k increase, mean fitness at $\hat{d} > 0$ increases.
 831 This results from the population being able to save costs of expensive learn-
 832 ing, while still producing quality social information, by allocating necessary
 833 individual learning to when it is really needed, when $t = 0$. A major oppos-
 834 ing force is the rate of change u , which reduces mean fitness at $\hat{d} > 0$. As u
 835 increases, the population spends less and less time at $t > 0$, and so reaps less
 836 benefit from any improvements in social information Q . Finally, the slower
 837 f increases with d , the higher mean fitness at $\hat{d} > 0$.

838 REFERENCES

- 839 Altenberg, L. 2012. Resolvent positive linear operator exhibit the reduc-
 840 tion phenomenon. *Proceedings of the National Academy of Sciences USA*
 841 109:3705–3710.
- 842 Altenberg, L. in press. The evolution of dispersal in random environments
 843 and the principle of partial control. *Ecological Monographs*.
- 844 Baum, W. M., P. J. Richerson, C. M. Efferson, and B. M. Paciotti. 2004. Cul-
 845 tural evolution in laboratory microsocieties including traditions of rule
 846 giving and rule following. *Evolution and Human Behavior* 25:305–326.
- 847 Boyd, R., and P. J. Richerson. 1995. Why does culture increase human adapt-
 848 ability? *Ethology and Sociobiology* 16:125–143.
- 849 Boyd, R., and P. Richerson. 1996. Why culture is common, but cultural evo-
 850 lution is rare. *Proceedings of the British Academy* 88:77–93.
- 851 Boyd, R., and P. J. Richerson. 1985. *Culture and the Evolutionary Process*.
 852 Chicago: Univ Chicago Press.
- 853 Cavalli-Sforza, L. L., and M. W. Feldman. 1981. *Cultural transmission and*
 854 *evolution: a quantitative approach*. Princeton: Princeton University Press.
- 855 Cohen, D. 1966. Optimizing reproduction in a randomly varying environ-
 856 ment. *Journal of Theoretical Biology* 12:119–129.
- 857 Cohen, D. 1967. Optimizing reproduction in a randomly varying environ-
 858 ment when a correlation may exist between the conditions at the time a
 859 choice has to be made and the subsequent outcome. *Journal of Theoretical*
 860 *Biology* 16:1–14.
- 861 DeWitt, T. J., A. Sih, and D. S. Wilson. 1998. Costs and limits of phenotypic
 862 plasticity. *Trends in Evolution and Ecology* 13:77–81.
- 863 Efferson, C., P. J. Richerson, R. McElreath, M. Lubell, E. Edsten, T. M. War-
 864 ing, B. Paciotti, and W. Baum. 2007. Learning, Productivity, Noise: An

- 865 Experimental Study of Cultural Transmission on the Bolivian Altiplano.
866 *Evolution and Human Behavior* 28:11–17.
- 867 Efferson, C., R. Lalive, P. J. Richerson, R. McElreath, and M. Lubell. 2008.
868 Conformists and mavericks: The empirics of frequency-dependent cul-
869 tural transmission. *Evolution and Human Behavior* 29:56–64.
- 870 Enquist, M., K. Eriksson, and S. Ghirlanda. 2007. Critical Social Learning: A
871 Solution to Rogers's Paradox of Nonadaptive Culture. *American Anthro-
872 pologist* 109:727–734.
- 873 Eriksson, K., and J. C. Coultas. 2009. Are people really conformist-biased?
874 An empirical test and a new mathematical model. *Journal of Evolutionary
875 Psychology* 7:5–21.
- 876 Eriksson, K., M. Enquist, and S. Ghirlanda. 2007. Critical points in current
877 theory of conformist social learning. *Journal of Evolutionary Psychology*
878 5:67–87.
- 879 Eriksson, K., and P. Strimling. 2009. Biases for acquiring information
880 individually rather than socially. *Journal of Evolutionary Psychology*
881 7:309–329.
- 882 Fox, S. 2002. “These are things that are really happening”: Inuit perspectives
883 on the evidence and impacts of climate change in Nunavut. *Pages 12–53
884 of: Krupnik, I., and D. Jolly (eds), The Earth is Faster Now: Indigenous
885 Observations of Arctic Environmental Change*. Fairbanks, Alaska: Arctic
886 Research Consortium of the United States.
- 887 Getty, T. 1996. The maintenance of phenotypic plasticity as a signal detection
888 problem. *The American Naturalist* 148:378–385.
- 889 Green, D. M., and J. A. Swets. 1966. *Signal detection theory and psychophysics*.
890 New York, NY: John Wiley and Sons Inc.
- 891 Hedrick, P. W., M. E. Ginevan, and E. P. Ewing. 1976. Genetic polymorphism
892 in heterogenous environments. *Annual Review Ecology Systematics* 7:1–32.
- 893 Henrich, J. 2001. Cultural Transmission and the Diffusion of Innovations:
894 Adoption dynamics indicate that biased cultural transmission is the pre-
895 dominate force in behavioral change and much of sociocultural evolution.
896 *American Anthropologist* 103:992–1013.
- 897 Henrich, J., and R. Boyd. 1998. The evolution of conformist transmission and
898 between-group differences. *Evolution and Human Behavior* 19:215–24.
- 899 Henrich, J., S. Heine, and A. Norenzayan. 2010. The Weirdest People in the
900 World? *Behavioral and Brain Sciences*.
- 901 Jacobs, R. C., and D. T. Campbell. 1961. The perpetuation of an arbitrary
902 tradition through several generations of a laboratory microculture. *The
903 Journal of Abnormal and Social Psychology* 62:649–658.
- 904 Kameda, T., and D. Nakanishi. 2002. Cost-benefit analysis of social/cultural
905 learning in a non-stationary uncertain environment: An evolutionary

- 906 simulation and an experiment with human subjects. *Evolution and Hu-*
907 *man Behavior* 23:373–393.
- 908 Levins, R. 1968. *Evolution in Changing Environments*. Princeton: Princeton
909 University Press.
- 910 McElreath, R. 2004. Social learning and the maintenance of cultural vari-
911 ation: An evolutionary model and data from East Africa. *American An-*
912 *thropologist* 106:308–321.
- 913 McElreath, R. 2010. The coevolution of genes, innovation and culture in
914 human evolution. *Pages 451–474 of: Silk, J., and P. Kappeler (eds), Mind*
915 *The Gap: Tracing the Origins of Human Universals*. Springer.
- 916 McElreath, R., and P. Strimling. 2008. When natural selection favors learning
917 from parents. *Current Anthropology* 49:307–316.
- 918 McElreath, R., M. Lubell, P. J. Richerson, T. Waring, W. Baum, E. Ed-
919 sten, C. Efferson, and B. Paciotti. 2005. Applying evolutionary models
920 to the laboratory study of social learning. *Evolution and Human Behavior*
921 26:483–508.
- 922 McElreath, R., A. V. Bell, C. Efferson, M. Lubell, P. J. Richerson, and T. War-
923 ing. 2008. Beyond existence and aiming outside the laboratory: Estimating
924 frequency-dependent and payoff-biased social learning strategies. *Philo-*
925 *sophical Transactions of the Royal Society B* 363:3515–3528.
- 926 McElreath, R., A. Wallin, and B. Fasolo. in press. The Evolutionary Rational-
927 ity of Social Learning. *In: Hertwig, R., U. Hoffrage, and the ABC Research*
928 *Group (eds), Simple Heuristics in a Social World*. Oxford University Press.
- 929 McNamara, J. M., and S. R. X. Dall. 2011. The evolution of unconditional
930 strategies via the ‘multiplier effect’. *Ecology Letters* 14:237–243.
- 931 Morgan, T. J. H., L. E. Rendell, M. Ehn, W. Hoppitt, and K. N. Laland. 2012.
932 The evolutionary basis of human social learning. *Proceedings of the Royal*
933 *Society London B* 279:653–662.
- 934 Nakahashi, W., J. Wakano, and J. Henrich. in press. Adaptive social learn-
935 ing strategies in temporally and spatially varying environments. *Human*
936 *Nature*.
- 937 Perreault, C., C. Moya, and R. Boyd. in press. A Bayesian approach to the
938 evolution of social learning. *Evolution and Human Behavior*.
- 939 Rendell, L., R. Boyd, D. Cownden, M. Enquist, K. Eriksson, M. W. Feldman,
940 L. Fogarty, S. Ghirlanda, T. Lillicrap, and K. N. Laland. 2010. Why copy
941 others? Insights from the social learning strategies tournament. *Science*
942 328:208–213.
- 943 Rogers, A. R. 1988. Does biology constrain culture? *American Anthropologist*
944 90:819–831.
- 945 Schreiber, S. J. 2010. Interactive effects of temporal correlations, spatial het-
946 erogeneity and dispersal on population persistence. *Proceedings of the*

- 947 *Royal Society London B* 277:1907–1914.
- 948 Sih, A., M. C. O. Ferrari, and D. J. Harris. 2011. Evolution and behavioural
949 responses to human-induced rapid environmental change. *Evolutionary*
950 *Applications* 4:367–387.
- 951 Tufto, J. 2000. The evolution of plasticity and nonplastic spatial and tem-
952 poral adaptations in the presence of imperfect environmental cues. *The*
953 *American Naturalist* 156:121–130.
- 954 Wakano, J. Y., and K. Aoki. 2007. Do social learning and conformist bias
955 coevolve? Henrich and Boyd revisited. *Theoretical Population Biology*
956 72:504–512.
- 957 Weatherhead, E., S. Gearheard, and R. G. Barry. 2010. Changes in weather
958 persistence: Insight from Inuit knowledge. *Global Environmental Change*
959 3:523–528.
- 960 Williams, P. D., and A. Hastings. 2011. Paradoxical persistence through
961 mixed-system dynamics: towards a unified perspective of reversal be-
962 haviours in evolutionary ecology. *Proceedings of the Royal Society London*
963 *B* 278:1281–1290.