

ROGERS' PARADOX RECAST AND RESOLVED: POPULATION STRUCTURE AND THE EVOLUTION OF SOCIAL LEARNING STRATEGIES

Luke Rendell,^{1,2} Laurel Fogarty,¹ and Kevin N. Laland¹

¹*School of Biology, University of St Andrews, Bute Building, Westburn Lane, St Andrews, Fife KY16 9TS, United Kingdom*

²*E-mail: ler4@st-andrews.ac.uk*

Received February 11, 2009

Accepted July 25, 2009

We explore the evolution of reliance on social and asocial learning using a spatially explicit stochastic model. Our analysis considers the relative merits of four evolved strategies, two pure strategies (asocial and social learning) and two conditional strategies (the "critical social learner," which learns asocially only when copying fails, and the "conditional social learner," which copies only when asocial learning fails). We find that spatial structure generates outcomes that do not always conform to the finding of earlier theoretical analyses that social learning does not enhance average individual fitness at equilibrium (Rogers' paradox). Although we describe circumstances under which the strategy of pure social learning increases the average fitness of individuals, we find that spatial structure introduces a new paradox, which is that social learning can spread even when it decreases the average fitness of individuals below that of asocial learners. We also show that the critical social learner and conditional social learner both provide solutions to the aforementioned paradoxes, although we find some conditions in which pure (random) social learning out-competes both conditional strategies. Finally, we consider the relative merits of critical and conditional social learning under various conditions.

KEY WORDS: Culture, evolution, learning, social learning strategy, spatial model.

Social learning is widespread in animals, and underpins both the cultural diversity and the considerable ecological success of humans. Although there has been considerable and recent theoretical analysis of the evolution of reliance on social learning, understanding of the relative effectiveness of different social learning strategies remains limited. Here, the term social learning refers to learning through observation of, or interaction with, another animal (Heyes 1994), which can occur by a number of psychological mechanisms including enhancement effects, imitation, and emulation (Whiten and Ham 1992; Heyes 1994; Hoppitt and Laland 2008). Conversely, asocial or individual learning refers to learning that occurs independently of any social influence. Historically, scientists have tended to assume that animals should rely on so-

cial learning whenever they can, but recent mathematical analyses reveal that this is incorrect, and that some mixture of social and asocial learning is expected to occur in a changing environment (Boyd and Richerson 1985; Rogers 1988; Feldman et al. 1996; Henrich and McElreath 2003; Enquist et al. 2007). This finding resembles the producer–scrounger frequency dependence observed in social foraging models (Barnard and Sibly 1981; Giraldeau and Caraco 2000). Asocial learners are information producers, who typically incur additional costs associated with learning from a direct interaction with the environment, whereas social learners are information scroungers, who can obtain information relatively cheaply from others, but are vulnerable to acquiring outdated or inappropriate information this way in a changing environment

(Kameda and Nakanishi 2002). Consequently, theoretical studies predict a mixture of social and asocial learning at equilibrium.

Anthropologist Alan Rogers (1988) first pointed out the “paradox” inherent in the observation that the fitness of social learners at this polymorphic equilibrium would be no greater than the average individual fitness in a population of asocial learners. When rare, the fitness of social learners exceeds that of asocial learners, but declines with frequency as there are fewer asocial learners producing adaptive information in a changing environment. The population evolves to a mixed evolutionarily stable strategy (ESS) where, by definition, the fitness of social learners equals that of asocial learners (Giraldeau et al. 2003; Henrich and McElreath 2003). This finding is now commonly known as Rogers’ paradox (Boyd and Richerson 1985), so called because it contrasts with a commonly held assertion that culture enhances fitness. Although Rogers’ result is not inherently paradoxical, it appears to conflict with the observation that social learning underlies the effect of human culture on our ecological success and population growth.

One resolution to this conundrum is to recognize that, in a changing environment, selection ought to have fashioned in our minds specific evolved rules (Boyd and Richerson 1985), or “social learning strategies” (Laland 2004), that specify the circumstances under which individuals should exploit information from others, and from whom they should learn. Previous theoretical studies have established that the average individual fitness at equilibrium can be enhanced if individuals switch between reliance on asocial and social learning (Boyd and Richerson 1995; Kameda and Nakanishi 2003). For instance, Enquist et al. (2007) showed that a strategy of “critical social learning,” where individuals only adopt asocial learning if social learning proves unsatisfactory, outcompetes pure asocial and social learning strategies. Boyd and Richerson (1985) also showed that average fitness is higher than that in a population of asocial learners if social learners can improve their learned behavior so that there is cumulative cultural evolution.

In this study, we extend the above analyses in three important respects. First, we investigate the effects of social and asocial learning in a spatially explicit context. There are several reasons why a spatial framework is appropriate for such analyses. Studies in other contexts have shown that spatial factors can profoundly affect evolutionary outcomes (Nowak and May 1992; Kerr et al. 2006; Silver and Di Paolo 2006). Some human cultural phenomena, such as agricultural practices, are physically grounded in space (Durham 1991). Moreover, social learning is now known to be widespread in vertebrates and even some invertebrates (Heyes and Galef Jr. 1996; Leadbeater and Chittka 2007), many of which are sedentary and/or territorial. In such cases, an analysis that recognizes that individuals are often more likely to learn from their near neighbors is appropriate. More generally, by compar-

ing well-mixed and spatially structured populations, our analyses allow us to characterize the extremities of a range of unstructured to structured populations. Learning in a structured population is a special case of bias in social learning, where nearby individuals are preferred as models to distant ones, and in this respect spatially explicit analyses are more generally instructive with respect to the effects of bias (Kameda and Nakanishi 2002). Below we show that spatial structure strongly affects the outcome of the aforementioned analyses into the relative merits of social and asocial learning.

Second, previous analyses have been generally reliant on deterministic models (in the sense that they ignore random events), although one recent exception is Whitehead and Richerson (2009). However, stochasticity could play an important role in affecting the balance of social and asocial learning, and the nature of the equilibria reached. Our analysis uses a stochastic model that allows individuals to disperse and to learn either locally or globally, thereby allowing us to tease apart the effects of drift and space.

Third, and finally, our analysis investigates the impact of a number of parameters that potentially affect reliance on social and asocial sources of information, but hitherto have either not been explored, or not been investigated in combination with each other. These parameters include the degree of temporal and spatial variation in the environment, environmental harshness (the extent to which suboptimal behavior reaps fitness benefits), the number of environmental states (which equates to the number of different ways of performing a suboptimal behavior), and the relative costs of social and asocial learning.

We show that spatial structure reinforces Rogers’ paradox, because social learning can spread even when it decreases the average fitness of individuals below that of asocial learners. We also show that there are circumstances under which the strategy of pure, unbiased, social learning increases the average fitness of individuals above that of asocial learners (unbiased here means choosing a model to copy at random). We find that two conditional strategies, the critical social learner (learn asocially only when copying fails) and conditional social learner (copy only when asocial learning fails), can both provide solutions to the aforementioned paradoxes, although we also find some conditions in which pure social learning outcompetes both conditional strategies, and so where the paradox remains. Finally, we consider the relative merits of critical and conditional social learning across a range of conditions. In the latter respect, we extend the findings of Enquist et al. (2007) to a stochastic, spatially explicit framework, and across a broader set of conditions.

The Model

We specify a square, $x \times x$, toroidal environment, where each cell is occupied by a single individual, such that $n = x^2$ is the

population size. Each cell has an associated environmental state, which can change over time, and environmental states can vary in space (i.e., between cells). In each time step, all individuals exhibit a behavior (phenotype), which can change over time, according to the dictates of an evolved learning strategy (their genotype), which is fixed, and governs whether and when individuals learn asocially or socially. The neighborhood of an individual consists of the eight surrounding cells (the Moore neighborhood). For mathematical convenience, we assume haploid asexual reproduction, such that individuals have only one parent.

ENVIRONMENTAL VARIABILITY

Each cell has a value representing its environmental state s , which takes an integer value between 1 and N_s , with adjacent integers representing similar environmental states. We implemented a ring structure making it possible for environments to step from N_s to 1, and vice versa, such that there were no environmental “end states.” We considered two types of environmental variation: (1) temporal only, where all cells have the same environmental state (s value) but this value can change over time, and (2) spatio-temporal, where, in addition, different cells can possess different s values. In the temporal-only condition, all cells switch simultaneously to a new, randomly chosen, value in the range $[1, N_s]$ with probability p_s . In the spatio-temporal condition, environmental variation occurs in two ways—perturbation of patches and single-cell changes. These two types of variation allowed us to control the level of spatial autocorrelation while also maintaining a reasonable degree of ecological validity. Perturbation events convert all the cells in a randomly placed square to a single environmental state in the range $[1, N_s]$, with each state having equal likelihood, and with a single perturbation event in each iteration (we chose a single event per iteration for computational convenience; increasing the number of perturbations would have the simple effect of increasing the rate of environmental change, a parameter that in any case we varied systematically in our simulations). The size of the perturbed square followed a power law distribution with the side length given by $8R^{-1/6}$ where R is uniform random in the interval $[0, 1]$. This perturbation regime was chosen for ecological validity as it produces relatively many small perturbations and occasionally large ones, as observed in real-world ecosystems (Langmead and Sheppard 2004). The use of a toroidal environment ensured that all cells had an equal probability of being affected by a perturbation event. Each individual cell subsequently switches state by a single step (i.e., from state 3 to state 2 or 4) with probability p_s . We implemented two conditions for these single step changes: cells change either (1) toward the average of their neighbors, if that average is different from its current state (spatially correlated condition) or (2) in a random direction (random condition). These conditions produce variation in the level of spatial autocorrelation, measured as the probability

of a cell’s neighbor being in the same environmental state as itself (henceforth p_n). In the temporal-only condition, $p_n = 1$, while in the spatio-temporal condition, spatially correlated change gives higher values of p_n than random change.

BEHAVIOR, PAYOFF, AND FITNESS

Each individual is characterized by behavior, b , which can take any integer value in the same range as the environment, 1 to N_s , and can change in each model iteration. Each individual’s genotype specifies a learning strategy. We consider two pure strategies, a for asocial learners and s for random social learners, as well as two conditional strategies, conditional social learner and critical social learner (described above). An individual’s fitness, W , is defined by the difference between the environmental state (s) in the cell it occupies and the behavior it currently shows, minus the cost of the learning strategy it uses

$$W = h^{-S} - c_{\text{strat}}, \quad (1)$$

where S is the number of steps between b and s implemented as a ring,

$$S_{(b,s)} = \begin{cases} |s - b| & \text{if } |s - b| \leq \lfloor N_s/2 \rfloor \\ N_s - |s - b| & \text{if } |s - b| > \lfloor N_s/2 \rfloor \end{cases},$$

h^{-S} is the payoff associated with being S steps away from the ideal behavior, and c_{strat} is the cost of the learning strategy (c_a or c_s for the costs of asocial or social learning, respectively, and some additive combination of these in the case of the two conditional strategies). In the case that $c_{\text{strat}} > h^{-S}$, W is given a lower bound of zero. The parameter h can be interpreted as the “harshness” of the environment, as increasing values increment the fitness penalty of behavior not matching the environment. Higher h values also result in an increasingly nonlinear payoff function, whereas as $h \rightarrow 1$ the payoff function approaches a linear, or “risk-neutral” form. Note that this incremental fitness structure differs from previous models (e.g., Rogers 1988; Enquist et al. 2007), which assumed a fitness payoff for learning the correct behavior only, with any other behavior receiving no payoff. In each model iteration, each individual reproduces with probability W , and the resultant offspring always replaces an existing individual such that the population size remains constant across iterations (see Dispersal below).

MUTATION

In general, offspring carry the same learning strategy as their parent, but at every reproduction event there is a fixed probability, μ , of mutation, which results in an offspring with an alternative strategy. This mutation is how new strategies are introduced.

DISPERSAL

We assume a juvenile developmental period during which offspring acquire their parent’s behavior prior to dispersal. This

occurs by asocial learning in the same environment as the parent in the case of asocial learners, and by observing the parent in the case of social learners. We simulate two dispersal conditions. Local dispersal results in offspring replacing a randomly chosen neighbor of the parent. In contrast, global dispersal results in offspring replacing an individual chosen at random from the entire habitat, excluding the parent. In both conditions, individuals are chosen to be replaced irrespective of their fitness and age. Reproduction is synchronous, such that it is not possible for offspring to be replaced in the iteration in which they are born.

LEARNING

We assume asocial learners sample their environment and learn the appropriate behavior, such that $S = 0$, with probability p_{iOK} , and acquire a random behavior with probability $1 - p_{iOK}$. Asocial learners pay a fixed cost c_a irrespective of the learning outcome, such that if $p_{iOK} = 1$ then their fitness is constant at $1 - c_a$. Social learners, in contrast, match the behavior of a randomly chosen demonstrator individual at each iteration, with fixed cost c_s , with demonstrators either chosen from the neighborhood (local copying) or the entire population (global copying). Consistent with earlier theory (Boyd and Richerson 1985; Rogers 1988) we assume $c_a \gg c_s$. Initially, we restrict our analysis to these two pure strategies and set $p_{iOK} = 1$, but subsequently introduce the two conditional strategies, critical social learner and conditional social learner (Enquist et al. 2007). Critical social learners first use social learning, at cost c_s , but if this fails to produce a match to the environment, they then switch to asocial learning, at cost c_a . Conditional social learners first try asocial learning, at cost c_a , and only if this fails to produce a match, use social learning at cost c_s . As conditional social learning only differs from asocial learning if there is a chance that asocial learning will not work, following Enquist et al. (2007) we specify $0 < p_{iOK} < 1$ when considering the conditional strategies. Thus, with probability $1 - p_{iOK}$, asocial learners acquire a random behavior. We assume for these strategies that individuals can know a candidate behavior does not match the environment, without knowing exactly what the state of the environment is. Individuals using these conditional strategies are forced to accept the final learning outcome they receive. For example, if a critical social learner switches to asocial learning it will acquire a random behavior with probability p_{iOK} and is forced to accept that behavior even if it returns a lower fitness than the behavior it originally learned socially and rejected.

SPATIAL POPULATION STRUCTURE

We concentrated on two conditions, which we termed “local” and “global.” In the local condition, both dispersal and copying were local—offspring disperse only to the neighborhood and social learners copy only neighbors. In the global condition, dispersal and copying were with respect to the entire population, as de-

scribed above. These two conditions represent a maximum and minimum level of population structure within our model. Intermediate levels can be represented by simulations in which only learning is local (“local learning”) or only dispersal is local (“local dispersal”) and we considered both these cases in our initial exploration of spatial effects, while noting that the latter condition (local dispersal with global learning) is biologically rather less plausible than the other three. In this way, we could explore the effect of population structure without altering other conditions.

DEMOGRAPHY

In our general model, increasing or decreasing average individual fitness affects only the rate of population turnover, as we specify a fixed population size. Although this assumption has utility in helping to understand the relative efficacy of alternative learning strategies, it does not allow us to investigate the demographic consequences of changes in individual fitness. We therefore explored a variant of our model that did allow for demographic effects, by making the simple assumption that the probability of any individual surviving to the next iteration (p_{survive}) is related to their fitness by the equation $p_{\text{survive}} = p_{\text{surviveMin}} + (1 - p_{\text{surviveMin}})W$, where $p_{\text{surviveMin}}$ is a parameter that defines the minimum probability an individual survives a model iteration and thus sets the strength of viability selection in the model. Thus, it is possible for the population to vary in size between zero (extinction) and the maximum n , the latter representing the carrying capacity of the environment. We assume empty cells provide no model for social learners to copy, but are no more likely to be occupied by new offspring than cells that are already occupied.

SIMULATION DETAILS

Each model iteration consisted of the following sequence of events: learning, followed by reproduction and dispersal, and then environmental change. For most simulations, we set $x = 80$, giving a population size, n , of 6400 individuals. This value of n was chosen so as to be broadly representative of human populations (6400 is close to the median for horticultural and herding societies, and intermediate between hunter-gatherer and agrarian societies, Lenski 1974). We set $\mu = 0.0008$, equating to five individuals per complete generation when $n = 6400$. Although this is obviously high relative to natural rates of mutation in eukaryotes, we found that reducing this rate does not qualitatively affect our outcomes, and the higher rate offers significant computational advantages in terms of time to equilibrium. All individuals were behaviorally naïve at the start of each simulation (i.e., $b = 0$). Below we discuss the results of sensitivity analyses investigating the extent to which our conclusions are robust across a range of biologically plausible values of n and μ .

We explored the effects of spatial variation in environmental conditions by producing three spatial autocorrelation conditions,

where the autocorrelation is expressed as the probability (p_n) that two randomly chosen but neighboring cells have the same environmental state (which is effectively the autocorrelation but with a “spatial lag” of one cell). The three conditions were $p_n = 1$ (no spatial variation at all), $p_n \approx 0.44$, and $p_n \approx 0.2$. The latter two values are approximate as they result from stochastic simulation runs and so the actual realized value in each simulation varied around these values; the actual values were recorded during simulation runs. We ran these simulations with $N_s = 10$, and local learning and dispersal, because we wished to explore conditions in which Rogers’ paradox may not exist. We explored the effect of environmental harshness (h) in the range $1.1 \leq h \leq 5$.

Simulations were run for 2000 iterations, with 20 replications per parameter value, and from each simulation we recorded the mean proportion of each strategy over the final 250 iterations of the run. Sensitivity analyses established that either increasing the number of iterations over which each simulation was run to 10,000, or increasing the number of runs per parameter set to 200, produced no change in the results.

We first compared the performance of a pure unbiased social learning strategy invading a population of pure asocial learners and, unless otherwise specified, assumed $p_s = 0.1$, $c_s = 0$, $p_{iOK} = 1$, and c_a increasing from 0.01 to 0.7 in steps of 0.01. We went on to investigate the evolution of conditional strategies, again starting from a population of asocial learners. We examined all four genotypes (the two pure and two conditional strategies) simultaneously, by considering the dynamics of a population of asocial learners capable of mutating to the other three genotypes. As conditional social learning only makes sense when there is a possibility of individual learning not working (i.e., $p_{iOK} < 1$), we made this comparison with $p_{iOK} = 0.5$. For these analyses, we introduced spatial variation in the environment by setting $p_n \approx 0.44$.

Results

ANALYSIS OF PURE STRATEGIES

Effect of local dispersal and learning

We first examined the performance of the pure asocial and social learning strategies in a spatially homogenous environment ($p_n = 1$). At the outset, we note that the global condition produced a qualitatively good fit to analytical expectations (Rogers 1988; Enquist et al. 2007) with respect to the magnitude of c_a at which social learners will invade, the resultant average final frequency of social learners, and the mean individual fitness in the population at run termination (Fig. 1). Similar to Rogers’s (1988) model, in the global condition the mean individual fitness approximates that expected in a population of entirely asocial learners, although it is marginally higher than expected for large c_a (Fig. 1B), whereas social learning frequencies were slightly

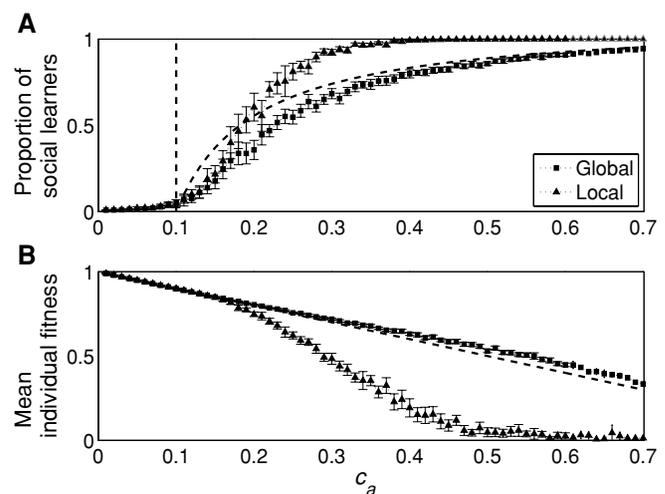


Figure 1. The effect of spatial population structure: global versus local dispersal and learning, showing (A) the proportion of social learners, and (B) mean individual fitness, at run termination, plotted against the cost of asocial learning (c_a). In both, $N_s = 1000$, $h = 2$, and $p_n = 1$. Points are means from 20 runs \pm 95% confidence intervals. The vertical dashed line in the upper panel shows where analytical models (Enquist et al. 2007) predict nonzero equilibrium levels of social learners, with the curved dashed line giving the analytical prediction for the frequency of the social learning strategy. The dashed line in the lower panel shows the expected mean individual fitness in a population of asocial learners only.

lower than the analytical prediction (Fig. 1A). These discrepancies can be attributed to minor differences in the assumptions of our model compared with the cited analytical treatments (see Discussion).

Local dispersal and learning typically generates higher average final frequencies of social learners than global dispersal and learning, particularly for larger values of c_a , and can lead to social learners reaching effective fixation (with an expected frequency of $1 - \mu$). Moreover, under local conditions the mean individual fitness in the population is predicted to be less than that expected for a population consisting of asocial learners, under conditions in which the frequency of social learners exceeds the analytical prediction. Thus, when population structure is imposed, not only does social learning invade and fail to increase mean individual fitness, it may actually invade to fixation and reduce fitness relative to a population containing only asocial learners.

The observation that social learning spreads to near fixation in spite of the fact that it reduces average individual fitness introduces a new dimension to the debate surrounding Rogers’ paradox. This observation can be explained by comparing the fitness of social learners that have at least one asocial learner in their neighborhood with those that do not (Fig. 2). When learning and dispersal are local, social learners with at least one asocial learner in their neighborhood have greater fitness than both social

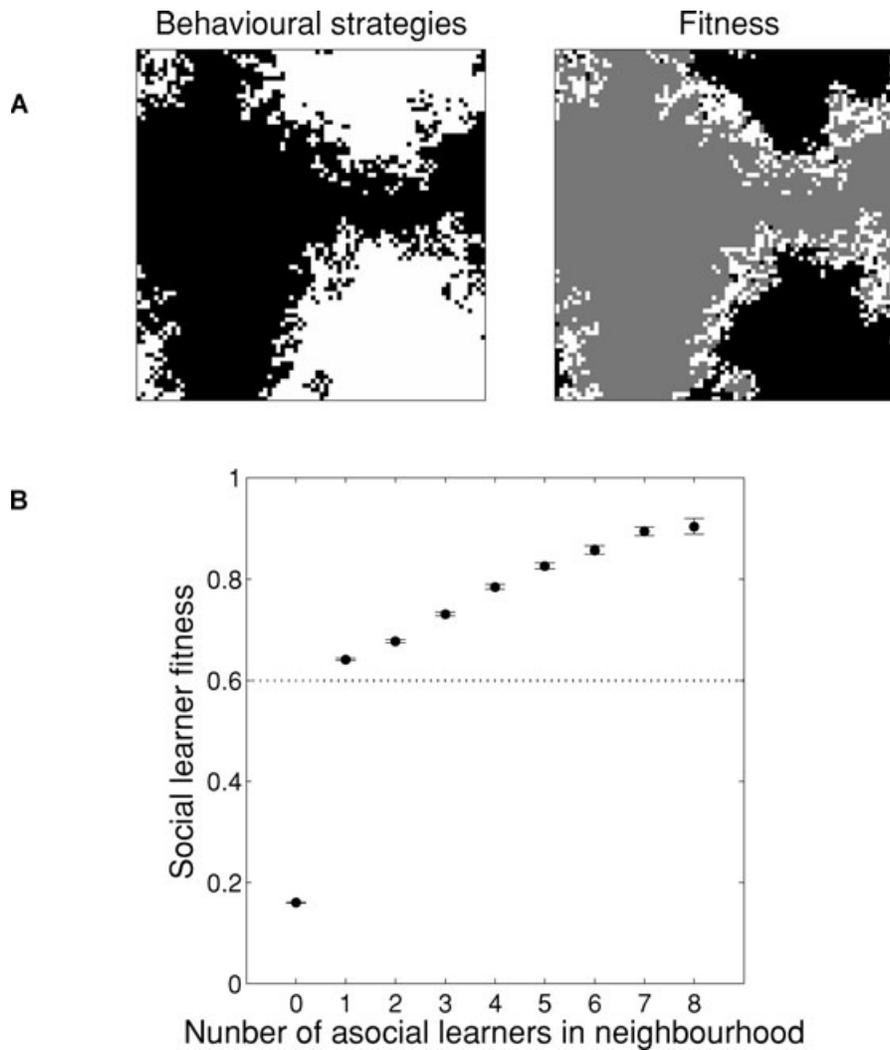


Figure 2. (A) Snapshot of a running simulation with local learning and dispersal and a spatially uniform environment, taken within a few iterations of an environmental state change. Left panel shows spatial distribution of learning strategies (black = asocial learner, white = social learner). Right panel shows the fitness of each individual; the lighter the cell, the higher the fitness value. Asocial learners have a fixed fitness ($1 - c_a$), which appears as gray. Social learners in the border regions have the highest fitness (appearing white in the right panel) even though most of the social learners have much lower fitness (appearing black). (B) Mean fitness of social learners plotted against the number of asocial learners in their neighborhood over the same simulation (error bars show 95% CI). Dashed line shows fitness of asocial learners.

learners with no asocial learner in their neighborhood (Fig. 2A) and asocial learners (Fig. 2B). This creates an edge effect in the contact zones between genotypes, and because social learners have greater fitness in these zones they can continue to increase in frequency even when, on average, social learners have lower fitness than asocial learners. Asocial learners cannot invade the regions dominated by social learners because their social learning neighbors will have a higher fitness immediately after an asocial learner mutates into their midst.

We conducted further simulations to separate out the effects of “local learning” and “local dispersal.” We found that local learning with global dispersal produced identical results to global learning and dispersal. In contrast, local dispersal with global

learning gave intermediate results, where some fixation of social learning was observed, but in a restricted parameter range compared to the global case.

Thus far, in our model, increasing or decreasing average individual fitness affects only the rate of population turnover, as we specify a fixed population size. Although helping to understand the relative efficacy of various learning strategies, this feature means it is not straightforward to understand the demographic consequences of changes in average individual fitness. We therefore ran the above simulations in a modified version of our model that did allow for demographic effects, by making the simple assumption that the probability of any individual surviving to the next iteration is related to their fitness. When we ran this model

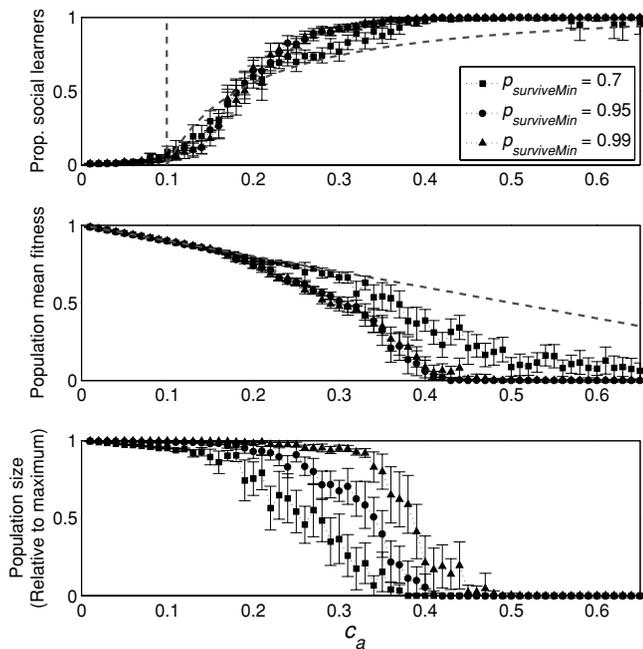


Figure 3. Demographic consequences of social learner invasion. Results are from simulations run under identical conditions to the local condition in Figure 1 with the exception that the probability an individual survives a model iteration is proportional to its fitness. Top panel shows the proportion of social learners, middle mean individual fitness, and bottom panel population size at run termination, plotted against the cost of asocial learning (c_a). Points are means from 20 runs \pm 95% confidence intervals, and the dashed lines in the top and middle panels are identical to those in Figure 1.

under the local conditions specified above, we obtained identical results in terms of social learner frequency and mean individual fitness, but with the confirmation that reduced mean individual fitness results in decreased population size, culminating in extinction when the average probability of reproduction approaches zero ($W \rightarrow 0$; Fig. 3). When viability selection was relatively weak (i.e., the minimum survival probability, $p_{\text{surviveMin}}$, was close to 1), we found a significant portion of the parameter space in which social learning would approach fixation, resulting in a stable population with reduced size compared to one containing only asocial learners. Stronger viability selection ($p_{\text{surviveMin}} \ll 1$) under local conditions does not prevent social learners spreading, in spite of the fact that they reduce mean individual fitness, to sufficiently high frequency that the population goes extinct (Fig. 3). These results did not change whether we started the simulations with populations size set to 1 (i.e., at carrying capacity) or 0.5 (half carrying capacity).

The above findings are generally robust across values of N_s (the number of possible environmental states) between 10 and 1000 and values of h (controlling the harshness of the penalty for suboptimal behavior) from 1.1 to 5, with one important exception.

For small N_s (e.g., $N_s = 10$), and in the local condition only, we found that for large c_a (e.g., $c_a \geq \sim 0.5$), the mean fitness of individuals in the population was elevated compared to that expected for a population of only asocial learners (Fig. 4A). Because under these conditions random copying does increase mean individual fitness, we have a set of conditions under which Rogers' original paradox does not exist. This fitness advantage of social over asocial learning is due to the assumption that offspring acquire the behavior of their parents during maturation—if offspring are assigned a random behavior at birth, high frequencies of social learners do not increase fitness (Fig. 4A). Note that the acquisition of parental behavior has no direct effect on offspring fitness, because they must undergo an independent learning round after dispersal before selection occurs. It does however increase the pool of potential models for social learners to copy. Because these models are offspring of individuals with relatively high fitness, on average, they will be performing locally adaptive behavior patterns, making them available for social learners in the neighborhood to copy. Thus, acquisition of parental behavior provides indirect fitness benefits to social learners. As these effects are reliant on the learning of locally adaptive behavior, low levels of spatial autocorrelation in the environment reduce the magnitude of this fitness increment (see below).

The simulations we report here all had the same mutation rate ($\mu = 0.0008$) and population size ($n = 6400$). We did explore the effects of varying these parameters, and the effects we report here are largely robust to such variation. Note that because μ is a rate, changing the population size (n) alters the absolute number of mutations per generation, and that therefore these two parameters are intimately related. Changes in μ or n have an intuitive effect on genotype frequencies: as expected, increasing mutation rate (or decreasing population size) increases the frequency of the less common genotype, and vice-versa. Accordingly, when social learning is common, mutation will tend to introduce asocial learners, and as the latter acquire adaptive behavior this means that the amount of correct information entering the population each generation covaries with the mutation rate. Increased mutation rates (or reduced population sizes) are therefore associated with increases in the mean fitness of individuals in the population. The reverse holds when the asocial learners dominate. In contrast, decreased mutation rates accentuate the fitness reductions associated with the local condition (e.g., Fig. 1), because new, correct, information enters the population less frequently.

Spatial variation in the environment

The level of spatial autocorrelation in the environment greatly affected the frequency of social learners in the population (Fig. 4A). Typically, spatial environmental variation reduces the benefits of social learning, because it reduces the probability that copied

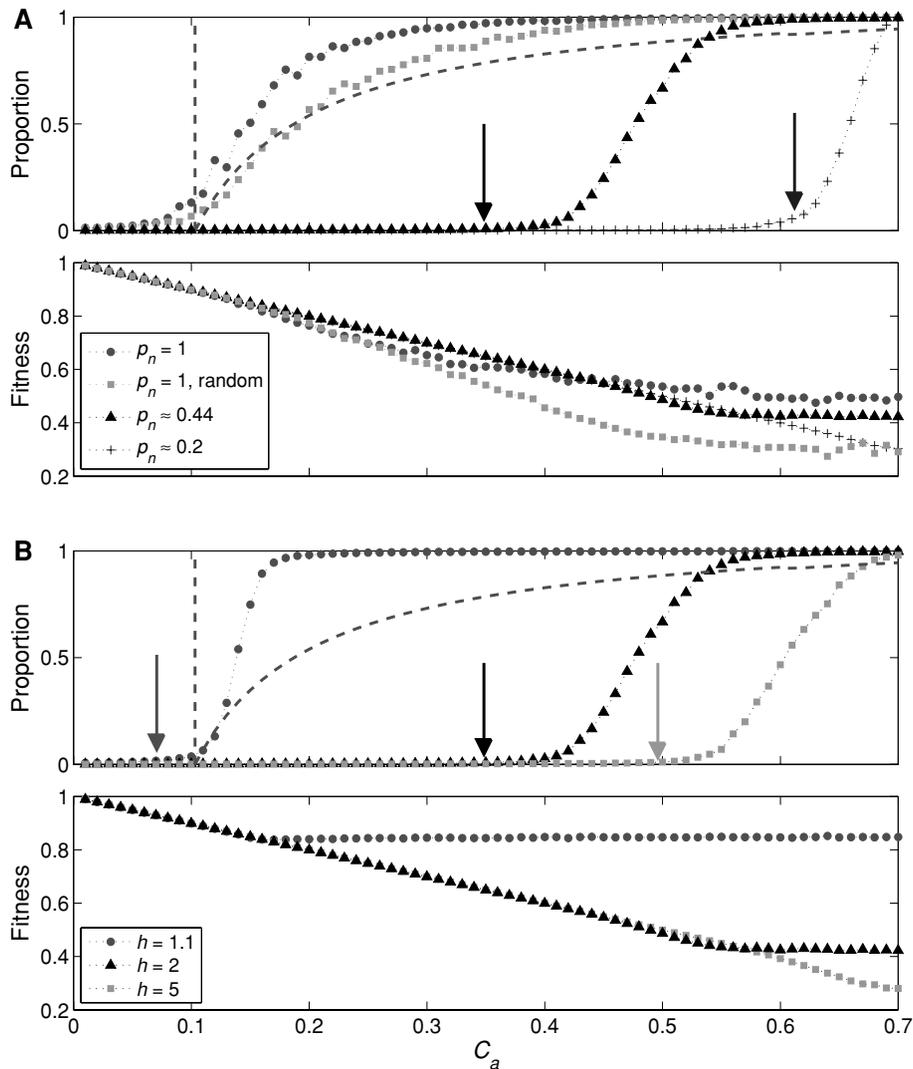


Figure 4. The effect of (A) spatial variation in the environment, and (B) the harshness of the fitness penalty for suboptimal behavior, on the proportion of social learners and mean individual fitness in simulated populations. Spatial variation is represented by the parameter p_n , the probability that two neighboring cells have the same environmental state, and harshness by the parameter h (see eq. 1). In all cases learning and dispersal are local, and if not otherwise specified $N_s = 10$, $h = 2$, and $p_n \approx 0.44$. Note that the data labeled “random” in (A) come from simulations in which there was no acquisition of parent’s behavior by offspring. Points are means from 20 runs, error bars are omitted for clarity. Dashed lines are as in Figure 1. Arrows indicate where an invading social learner has equal fitness to surrounding asocial learners according to the inequality specified in (3). In the case $p_n \approx 0.2$ in (A), note that when c_a is high (> 0.5), fitness, reproductive rate, and thus selection strength is low, such that drift effects are able to maintain low levels of social learners in the population even when their fitness is slightly lower than asocial learners.

individuals will have experienced the same environment as the copier. This means that social learning requires higher costs of asocial learning to invade a spatially variable environment, compared with a uniform one. When dispersal and learning are global, small values of p_n mean that social learning never becomes established. However, local learning partly negates these effects. In the local condition, at intermediate levels of spatial autocorrelation ($p_n \approx 0.44$), social learning can only become established at higher levels of c_a than the $p_n = 1$ case, although the subsequent transition to fixation occurs over a smaller range of c_a values.

This effect is exaggerated at $p_n \approx 0.2$, with social learners unable to invade until c_a is very large, because social learning has to be considerably cheaper than asocial learning if it is to have an advantage when copying an asocial learner directly only returns good information 20% of the time. Provided spatial autocorrelation is sufficiently high, local learning results in an increase in mean individual fitness above that expected for asocial learners, even when social learning reaches near fixation. This establishes further conditions for which Rogers’ paradox does not apply. Conversely, low spatial autocorrelation leaves mean individual fitness

at or near that expected for a purely asocial learning population, even with local learning.

In contrast to uniform environments, changes in environmental harshness (h) have major effects on the frequency of social learning in spatially varying environments (Fig. 4B). Increasing h has the effect of reducing the fitness payoff to social learners performing suboptimal behavior, the frequency of which increases with spatial variation. However small N_s values ameliorate the impact of increased h , because proportionally more behavior patterns will reap higher fitness dividends. This also means that lower values of h result in an escape from Rogers' paradox at relatively low levels of c_a , when N_s is low.

Predicting when social learning will invade

Here, we derive an analytical approximation for the conditions under which a pure social learning strategy will invade a population of asocial learners under local learning conditions. Near the asocial learning fixation boundary, the social learning genotype will increase when its expected fitness (W_s) exceeds $1 - c_a$, which is the constant fitness of asocial learners (assuming $p_{iOK} = 1$). In the invasion condition of a single social learner mutating into a neighborhood of asocial learners, the probability of that social learner copying correct behavior from a neighbor is given by the probability that the neighboring cell has the same environmental state multiplied by the probability that the neighboring cell has not changed state in the iteration because its occupant learned the correct behavior. We can write this as $p_n(1 - p_c)$, where p_c is the probability of a cell changing environmental state per iteration. When the environment is spatially uniform, p_c equals the model parameter p_s , and when it is not, $p_c \approx p_s + 80/n$ (the second term being the probability of a cell being included in a perturbation event of area $(8R^{-1/6})^2$ where the expectation of R is 0.5). When learning is correct, then the payoff is 1. When learning is not correct, with probability $1 - p_n(1 - p_c)$, then the expected payoff E_s approximates a weighted average of the payoffs when S (from eq. 1) is greater than zero, as it must be when environment and behavior do not match. Here, the weights are the probabilities of neighboring cells being a given number of environmental state steps apart. These considerations allow us to derive an expression for the expected fitness of an invading social learner (W_s) surrounded by asocial learners, or

$$\begin{aligned}
 W_s &= P(\text{learning is correct}) \\
 &\times 1 + P(\text{learning is incorrect}) \times E_s - c_s \quad (2) \\
 &= p_n(1 - p_c) + (1 - p_n(1 - p_c)) E_s - c_s.
 \end{aligned}$$

As successful invasion requires that $W_s > 1 - c_a$, the condition for invasion can be written as

$$c_a > 1 - p_n(1 - p_c) - (1 - p_n(1 - p_c)) E_s + c_s, \quad (3)$$

where

$$E_s = \frac{\sum_{i=1}^{\lfloor N_s/2 \rfloor} w_i h^{-i}}{\sum_{i=1}^{\lfloor N_s/2 \rfloor} w_i},$$

and w is a vector of the probabilities that two neighboring cells will be $[1, 2 \dots \lfloor N_s/2 \rfloor]$ environmental state steps apart. Because of the stochastic nature of our simulations and the complex spatial variation that resulted, we calculated w directly from the simulation runs. Note however that any behavior-payoff function can be used to calculate E_s for equation (2) provided it generates an expectation when behavior is suboptimal.

The predictions generated by this inequality conform reasonably well to our simulations (Fig. 4). Note that when there is no spatial variation ($p_n = 1$), no cost to social learning ($c_s = 0$), and an infinite sequence of possible environmental states ($N_s = \infty$), as posited by Rogers (1988) and Enquist et al. (2007), then $E_s = 0$, and (2) simplifies to $c_a > p_c$, as in their models. Inequality (3) specifies that whether a strategy of pure social learning will invade depends not only on the relative costs of social and asocial learning and the rate of environmental variation, but also on the level of spatial autocorrelation in the environment and the expected payoff of choosing a suboptimal behavior at random.

ANALYSIS OF CONDITIONAL STRATEGIES

Here, we consider the merits of two conditional strategies, conditional social learning (where asocial learning is attempted first followed by social learning if asocial learning fails) and critical social learning (where social learning is attempted first followed by asocial learning if a correct result is not obtained). We began our investigation by again comparing our model directly to the analytical results of Enquist et al. (2007), which requires $p_n = 1$. We found that critical social learning was able to both invade a population of pure asocial learners and increase individual mean fitness under the same conditions predicted by Enquist et al. (2007). Local learning and dispersal made very little difference to the pattern of results. We also examined all four genotypes (the two pure and two conditional strategies) simultaneously, by considering the dynamics of a population of asocial learners capable of mutating to any the other three genotypes. When $p_{iOK} < 1$, the expected fitness of asocial learners (W_a) is given by

$$W_a = p_{iOK} + (1 - p_{iOK}) E_a - c_a \quad (4)$$

where E_a is the expected payoff when asocial learning fails (i.e., the expected payoff of a randomly chosen behavior), given by

$$E_a = \frac{\sum_{i=1}^{N_s} h^{-S(1,i)}}{N_s}.$$

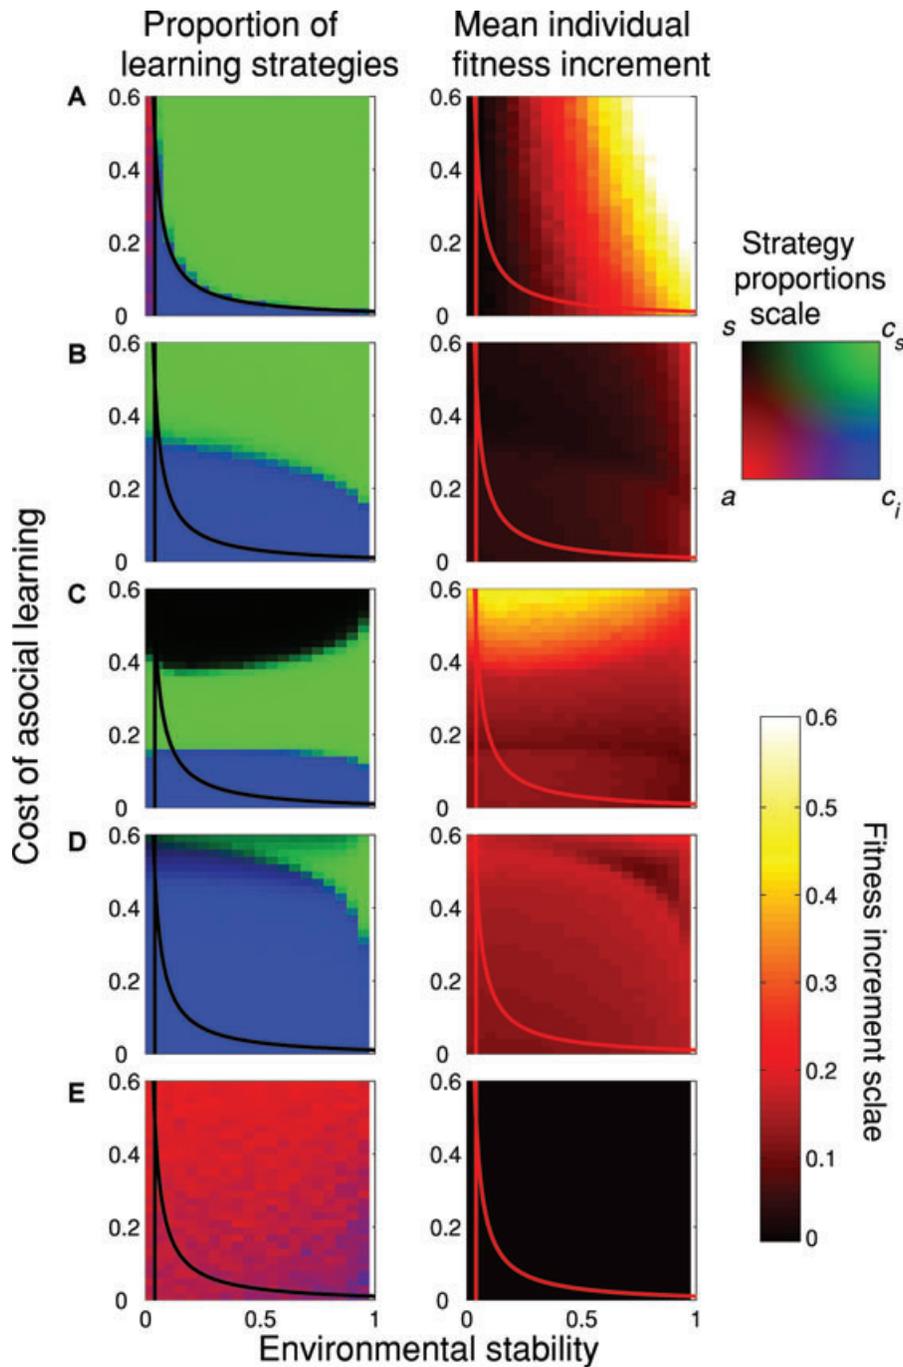


Figure 5. Frequencies of asocial (a , red), social (s , black), critical social (c_s , green), and conditional social (c_i , blue) learning strategies (left panel), and population fitness increment over that expected of a purely asocial learning population (right panel), with the cost of asocial learning plotted against environmental stability. Solid lines are estimates of the analytical predictions of Enquist et al. (2007) for a perfectly mixed population. Unless otherwise specified $N_s = 1000$, $h = 2$, $p_{iOK} = 0.5$, $c_s = 0.02$, and learning and dispersal are local. (A) No spatial variation, $p_n = 1$. (B) Effect of spatial environmental variation; $p_n \approx 0.44$. (C) Effect of reducing N_s ; $N_s = 10$ and $p_n \approx 0.44$. (D) Effect of reducing h ; $h = 1.1$ and $p_n \approx 0.44$. (E) Effect of global learning and dispersal, in a spatially variable environment, $p_n \approx 0.44$.

Our results are virtually identical to Enquist et al.'s (2007) consideration of three strategies (asocial plus the two conditional strategies) under similar conditions (Fig. 5A). Pure social learning never attained frequencies above the mutation rate in any of the si-

mulations, which is explained by the observation that social learners never had fitness higher than either asocial learning or either conditional strategy when surrounded by those strategies (Fig. 6). Results for the local and global conditions were comparable.

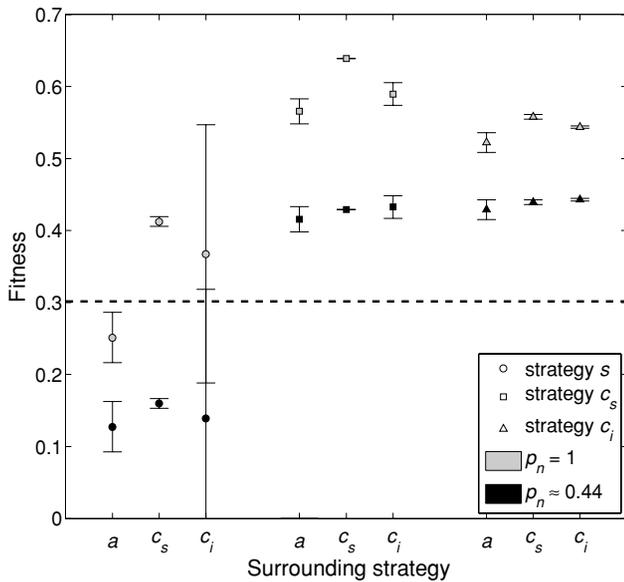


Figure 6. Mean fitness of social learners (*s*), critical social learners (*c_s*) and conditional social learners (*c_i*) when surrounded entirely by individuals of other strategies, in simulations under the same conditions as those in Figure 5A (spatially uniform environment, $P_n = 1$) and 5B (spatially heterogeneous environment, $P_n = 0.44$), with $c_a = 0.2$ and $p_c = 0.5$. Error bars show standard errors of the means. No data are shown for any strategy surrounded entirely by social learners, as pure social learners never attained sufficient frequency in these simulations to completely surround any cell. Dashed line shows fitness of asocial learners calculated using equation (4).

Conditional and critical social learning in spatially heterogeneous environments

However, our results differ markedly from earlier theory when we introduce spatial variation in the environment. In the local condition, we found that the parameter space where conditional social learning was favored over critical social learning to be significantly expanded, such that critical social learning only dominated at higher levels of c_a (Fig. 5B–D). In spatially heterogeneous environments, social learning is less effective than in uniform environments, because individuals are more likely to pick up inappropriate information from others. (This is illustrated by the fitness of social learners in Fig. 6). This has the effect of decreasing the probability that critical social learners will find an optimum behavior through social learning, which they attempt first, and consequently increases the likelihood that they will also pay the cost of asocial learning. Figure 6 illustrates how the fitness advantage of critical social learning over conditional social learners observed in a spatially homogeneous environment, is reversed in a spatially heterogeneous environment, because the comparative ineffectiveness of social learning means that critical social learners more frequently pay the cost of both forms of learning than do the conditional social learners.

In the global condition, in general, no strategy involving social learning reached high frequency, except in the region in which asocial learning is relatively cheap and the environment is highly stable (Fig. 5E). Under these conditions, conditional social learning can reach relatively high frequency, but even then, any increase in mean individual fitness is negligible. Surprisingly, this leads to the counter-intuitive prediction that a social learning strategy will be favored when asocial learning is cheap (see Discussion).

Once again, these qualitative findings are robust to variations in both the number of environmental states (N_s) and the harshness of the environment (h). We note, however, that where N_s is small and c_a is high, there are broad circumstances under which the pure strategy of random social learning is favored over both conditional strategies (e.g., Fig. 5C, where $N_s = 10$). Here the high probability of acquiring high fitness behavior by chance renders the additional cost of sometimes or always paying the cost of asocial learning unprofitable. Reducing the harshness of the environment increases the range of parameter space over which the conditional social learning strategy outcompetes the critical social learning strategy (Fig. 5D).

These results can be better understood with reference to expressions for the expected fitness of the conditional strategies, W_{crit} and W_{cond} for critical social learner and conditional social learner, respectively, given by

$$W_{crit} = p_{sOK} + (1 - p_{sOK})(p_{iOK} + (1 - p_{iOK})E_a) - (c_s + (1 - p_{sOK})c_a) \tag{5}$$

and

$$W_{cond} = p_{iOK} + (1 - p_{iOK})(p_{sOK} + (1 - p_{sOK})E_s) - (c_a + (1 - p_{iOK})c_s), \tag{6}$$

where the term p_{sOK} represents the probability that social learning returns the optimum behavior. Critical social learners outcompete conditional social learners when $W_{crit} > W_{cond}$, which simplifies to

$$E_a(1 - p_{iOK})(1 - p_{sOK}) + c_a p_{sOK} > E_s(1 - p_{iOK})(1 - p_{sOK}) + c_s p_{iOK}.$$

By denoting $E_s = E_a + \alpha$, we can simplify this expression further to give

$$c_a p_{sOK} - c_s p_{iOK} > \alpha(1 - p_{iOK})(1 - p_{sOK}). \tag{7}$$

This inequality illustrates how high values of the cost of asocial learning (c_a) and of the probability that social learning returns optimal behavior (p_{sOK}) will favor critical social learning. Environmental homogeneity (high p_n , low p_c) typically generates high values of p_{sOK} , because it makes it more likely that the copied individual will exhibit the correct behavior for the observer, giving

critical social learners the advantage. Conversely, environmental heterogeneity (low p_n , high p_c) reduces the efficacy of social learning (reduces p_{sOK}), allowing conditional social learners to invade. The inequality also shows how the greater the magnitude of the expected payoff from suboptimal copied behavior over randomly chosen behavior (α), the more likely that conditional social learning invades. Any such increase in α will be affected by environmental autocorrelation and also the parameter h .

Discussion

We have explored the evolution of different forms of learning using a spatially explicit stochastic model that incorporates a number of factors that have previously been absent or rare in theoretical literature on the evolution of social learning up until now. We began by validating our model by comparing it with earlier analytical treatments (Rogers 1988; Enquist et al. 2007) and demonstrated that it behaved comparably. There were nonetheless small, but notable, discrepancies in some cases. For example in the global condition shown in Figure 1, frequencies of social learning were lower, and mean individual fitness was slightly higher, than the analytical prediction. These minor differences can be attributed to two factors.

First, in the model of Enquist et al. (2007), cultural evolution (the spread of information by social learning) and genetic evolution (the change in strategy genotypes over time) were strictly separated, largely for reasons of analytical tractability, under the assumption that the former was rapid compared to the latter and thus likely to reach equilibrium before the next biological generation. In contrast, these processes occur concurrently in our models, such that selection could act against social learners to whom correct information had not yet diffused following environmental change, disfavoring social learners relative to the analytical case.

Second, our models assume that offspring acquire the same behavior as their parent, either directly because of vertical cultural transmission, or indirectly because through asocial learning each is exposed to the same environmental contingencies. Given that those parents with correct behavior are disproportionately more likely to reproduce, this factor would tend to increase mean individual fitness slightly. Although vertical transmission does not directly increment offspring fitness, it increases the pool of potential “demonstrators” performing high-fitness behavior patterns that are available to be copied, and hence provides peripheral fitness benefits to social learners. Despite this, the behavior of our models is qualitatively similar to those of Rogers (1988) and Enquist et al. (2007), which lends confidence that our extensions into analytically intractable contexts are well founded in existing theory.

Our most striking finding is that, when learning and dispersal are confined to local neighborhoods, social learning can invade to effective fixation and yet result in a population with lower mean individual fitness than would be expected with asocial learning alone. This can be explained by examining the fitness of different segments of the social learner population as the simulation runs. The most useful comparison is between those social learners that have at least one asocial learner in their neighborhood and those that do not. In the condition in which learning and dispersal are global, both these subpopulations have identical fitness. However, when learning and dispersal are local, social learners with at least one asocial learner in their neighborhood have greater fitness than both other social learners and those asocial learners that have at least one social learner in their neighborhood. In other words, there is an edge effect wherever social learners are in contact with asocial learners (social learners can acquire useful information from nearby asocial learners at low cost) and as dispersal is local, then the strategies can only replace each other in these edge zones. Because the social learning genotype has greater fitness in these contact zones, social learners increase in frequency even though they on average have lower fitness than asocial learners (Fig. 2). This perpetuates even though mutations can give rise to asocial learners within the social learner population because as soon as an asocial learner arises through mutation, the social learners surrounding it jump in fitness as they now have a useful source of information available, generating selection against the mutants. In contrast, social learners surrounded only by other social learners are isolated from such useful information, and become part of an information cascade relying on increasingly outdated information, with potentially detrimental consequences when the environment changes (Bikhchandani et al. 1992; Giraldeau et al. 2003). By this process, the ability of the population to track a changing environment is lost, so mean individual fitness falls almost to zero.

The fitness associated with differential learning strategies in our model can be considered in two ways. The first is to consider the fitness associated with learning to be just one component of an individual's fitness, and hence having a mean individual fitness of zero need not correspond to population shrinkage if we assume a baseline fitness level that maintains the population level. The second is to consider the fitness in our models as representing the entire fitness of the individuals concerned. Under this interpretation, the way our analyses artificially hold population size constant can make it difficult to interpret our results in terms of demographic consequences. The extension of our model to incorporate viability selection (Fig. 3) helps in this regard by showing that reductions in average individual fitness due to high levels of social learning in a population do indeed reduce population size, and can lead to extinction. Here, our results are very similar to those obtained by Whitehead and Richerson (2009). Although the

risks associated with over-reliance on social learning producing information cascades are well understood (Bikhchandani et al. 1992; Giraldeau et al. 2003), the role that population structure could play in producing cascade-like effects, and their population level evolutionary effects, has not previously been fully appreciated.

Although no real-world population would be as rigidly structured as those in our simulations, the local condition in our analysis represents one bound of the spectrum of possible structural constraints, with the equally unrealistic perfect mixing assumed by analytical models (and our global condition) representing the other. Real populations would be found somewhere between these two bounds, which means that any invasion of a pure strategy of social learning may reduce fitness. It is not hard to find examples of population structure with regard to information flow. One example is the Landga people of New Guinea, where “craftsmen report that, because of the great value of the skill [of stone adze construction], they will instruct only close relatives” (Stout 2002). Furthermore, spatial structure can be thought of as equivalent to any social learning bias that results in copying from only a select few behavioral models (Kameda and Nakanishi 2002). Thus, our analysis suggests that Rogers’ paradox is even stronger than originally thought; pure social learning does not just fail to increase fitness, it may even frequently reduce it.

We have also, however, described some conditions in which this paradox need not apply. Under the specified behavior and payoff structure, increasing the number of environmental states effectively increases the number of ways an individual can generate suboptimal solutions and thus have low relative fitness. The parameter h , which can be interpreted as representing the “harshness” of a particular environment, sets the size of the relative fitness increment (ΔW) available to organisms that perform the correct behavior compared with alternatives (as $h \rightarrow 1$, then $\Delta W \rightarrow 0$). Intuitively, it makes sense that as the fitness cost of acquiring a suboptimal behavior through copying decreases then that cost will become less significant relative to the cost of asocial learning, and so social learners will be favored. This pattern was observed in those simulations where h and N_s were manipulated. Under conditions when h and/or N_s are low and c_a sufficiently high, random social learning can increase the mean fitness of individuals. These conditions represent relatively benign ecological contexts, where almost any realistic behavior leads to some fitness benefit.

Our analysis of conditional strategies endorses the conclusions of Enquist et al. (2007) that the critical social learner and conditional social learner strategies also resolve Rogers’ paradox. Enquist et al. suggested that critical social learners outcompete conditional social learners over a broad area of the parameter space defined by the cost of asocial learning, c_a , and environmental stability ($1 - p_c$), and under the conditions they explored, our

models generate the same findings (Fig. 5A). However, the introduction of a spatially varying environment in our simulations significantly alters these conclusions (Fig. 5B–E) by revealing conditions under which the conditional social learning strategy is favored over the critical social learning strategy. Our analyses suggest that the cost of asocial learning is the more important factor in determining the switch between conditional social learning and critical social learning, with environmental stability having a relatively weak influence. The switch typically happens when the cost of asocial learning is around an order of magnitude higher than the cost of social learning. In circumstances in which social learning is effective, then critical social learning will be at an advantage, because it will tend to pay the cost of asocial learning less than conditional social learning. Conversely, if social learning is ineffective, then the reverse is true. Because increased environmental spatial variation decreases the effectiveness of social learning, it will broadly favor conditional social learning (see Fig. 6 and Inequality 7).

When learning and dispersal are global but the environment varies in space, it is very difficult to obtain an appropriate behavior through social learning. However, conditional social learning can invade when asocial learning is relatively unreliable, the environment is highly stable, and, counter intuitively (because prior theory has found that social learning is favored when the cost of asocial learning is high, e.g., Boyd and Richerson 1985), when the cost of asocial learning is low. Under these conditions, conditional social learners will have an advantage over asocial learners when asocial learning fails, and when there is a chance of learning something useful by social learning, which is higher when the environment is stable. However, the payoff advantage is small, and so will only have an effect when overall learning costs are low, hence the counterintuitive result of a social learning strategy invading a population of asocial learners when the cost of asocial learning is low. This is also the region where conditional social learners outperform critical social learners. The latter cannot become established in stable environments because it is difficult for the benefits of social learning to outweigh the low costs of asocial learning.

We also found that pure random social learning can outcompete either of the conditional genotypes when asocial learning was relatively unreliable, N_s is low and c_a is high (Fig. 5C). This was because a low number of possible environmental states means that even suboptimal payoffs from pure social learning can be better than more optimal payoffs associated with high learning costs. Here, the high probability of acquiring high fitness behavior means the additional burden of sometimes or always paying the cost of asocial learning carried by the conditional strategies renders them unprofitable. This finding suggests that random copying may be a good strategy in contexts in which it is difficult to produce something effective alone, but there is a range of viable

alternatives available to copy, all of which may produce something reasonably effective. Aesthetic craft production, such as pottery decoration or rug weaving, may fit these conditions, where the chances of producing something pleasing to the eye oneself may be relatively low, but there are a range of examples available to copy, all of which have previously been selected because they are to some degree, if not equally, pleasing to the eye. Recent evidence reveals that for such traits, neutral drift models provide a good fit. For example, Bentley et al. (2004) describe some specific examples of such cultural traits that appear to fit models of random drift that would be expected from random copying. Our analyses suggest that these contexts may favor adopting a random copying strategy.

Increasing environmental harshness (h) favors the critical social learning strategy, and reduced harshness the conditional social learning strategy (Fig. 5D; Inequality 7). We note that the formulation of Enquist et al. (2007), in which only one in an infinite number of behavior patterns confers any fitness benefit, represents an extreme in the possible distribution of h which favors the critical social learning strategy. The effect of harshness occurs because both conditional strategies deploy a combination of asocial and social learning, but when critical social learners learn asocially, and it fails to produce the correct solution, they are allotted a random behavior. In contrast, when conditional social learners learn asocially, and it fails to produce the correct solution, they receive the expected payoff of a social learner. Because offspring acquire their parent's behavior, then selection ensures that social learning is likely to acquire behavior closer, on average, to optimal than random. By magnifying the expected returns to a social learner performing a suboptimal behavior (increasing α in Inequality 7), reduced harshness benefits conditional social learners, because critical social learners switch to asocial learning if social learning returns suboptimal behavior, and therefore risk acquiring a random behavior rather than satisficing with "near misses" from social learning.

In summary, using a simple spatially explicit stochastic model, we have shown that spatial structure, including local learning and dispersal, can affect the evolution of social learning in ways that would be difficult to explore and predict using an analytical approach. The simulation framework we have presented here is potentially a useful vehicle for exploring more complex social learning strategies, such as "copy-in-proportion," or "copy-if-better" (Laland 2004), which would also be difficult to specify analytically.

ACKNOWLEDGMENTS

We are indebted to M. Enquist, K. Eriksson, and P. Strimling for a discussion of various aspects of this work. We thank A. Kandler, J. Kendal, M. Doebeli, and two anonymous reviewers for many helpful comments on the manuscript. This research was supported by the CULTAPTATION project (European Commission contract FP6-2004-NESTPATH-043434)

and by a grant from the Biotechnology and Biological Sciences Research Council (BB/C005430/1) to KNL.

LITERATURE CITED

- Barnard, C. J., and R. M. Sibly. 1981. Producers and scroungers—a general-model and its application to captive flocks of house sparrows. *Anim. Behav.* 29:543–550.
- Bentley, R. A., M. W. Hahn, and S. J. Shennan. 2004. Random drift and culture change. *Proc. R. Soc. Lond. B.* 271:1443–1450.
- Bikhchandani, S., D. Hirshleifer, and I. Welch. 1992. A theory of fads, fashion, custom, and cultural change as informational cascades. *J. Polit. Econ.* 100:992–1026.
- Boyd, R., and P. J. Richerson. 1985. *Culture and the evolutionary process*. Chicago Univ. Press, Chicago.
- . 1995. Why does culture increase human adaptability? *Ethol. Sociobiol.* 16:123–143.
- Durham, W. H. 1991. *Coevolution: genes, culture, and human diversity*. Stanford Univ. Press, Stanford, CA.
- Enquist, M., K. Eriksson, and S. Ghirlanda. 2007. Critical social learning: a solution to Rogers' paradox of non-adaptive culture. *Am. Anthropol.* 109:727–734.
- Feldman, M. W., K. Aoki, and J. Kumm. 1996. Individual versus social learning: evolutionary analysis in a fluctuating environment. *Anthropol. Sci.* 104:209–231.
- Giraldeau, L. A., and T. Caraco. 2000. *Social foraging theory*. Princeton Univ. Press, Princeton, NJ.
- Giraldeau, L.-A., T. J. Valone, and J. J. Templeton. 2003. Potential disadvantages of using socially acquired information. *Philos. Trans. R. Soc. Lond. B* 357:1559–1566.
- Henrich, J., and R. McElreath. 2003. The evolution of cultural evolution. *Evol. Anthropol.* 12:123–135.
- Heyes, C. M. 1994. Social learning in animals: categories and mechanisms. *Biol. Rev.* 69:207–231.
- Heyes, C., and B. G. Galef, Jr., eds. 1996. *Social learning in animals: the roots of culture*. Academic Press, San Diego, CA.
- Hoppitt, W., and K. N. Laland. 2008. Social processes influencing learning in animals: a review of the evidence. *Adv. Study Behav.* 38:105–165.
- Kameda, T., and D. Nakanishi. 2002. Cost–benefit analysis of social/cultural learning in a nonstationary uncertain environment: an evolutionary simulation and an experiment with human subjects. *Evol. Hum. Behav.* 23:373–393.
- . 2003. Does social/cultural learning increase human adaptability? Rogers's question revisited. *Evol. Hum. Behav.* 24:242–260.
- Kerr, B., C. Neuhauser, B. J. M. Bohannan, and A. M. Dean. 2006. Local migration promotes competitive restraint in a host–pathogen 'tragedy of the commons'. *Nature* 442:75–78.
- Laland, K. N. 2004. Social learning strategies. *Learn. Behav.* 32:4–14.
- Langmead, O., and C. Sheppard. 2004. Coral reef community dynamics and disturbance: a simulation model. *Ecol. Model.* 175:271–290.
- Leadbeater, E., and L. Chittka. 2007. Social learning in insects—from miniature brains to consensus building. *Curr. Biol.* 17:R703–R713.
- Lenski, G. 1974. *Human Societies*. McGraw-Hill, New York.
- Nowak, M. A., and R. M. May. 1992. Evolutionary games and spatial chaos. *Nature* 359:826–829.
- Rogers, A. 1988. Does biology constrain culture? *Am. Anthropol.* 90:819–813.
- Silver, M., and E. Di Paolo. 2006. Spatial effects favour the evolution of niche construction. *Theor. Popul. Biol.* 70:387–400.

Stout, D. 2002. Skill and cognition in stone tool production: an ethnographic case study from Irian Jaya. *Curr. Anthropol.* 43:693–722.

Whitehead, H., and P. J. Richerson. 2009. The evolution of conformist social learning can cause population collapse in realistically variable environments. *Evol. Hum. Behav.* 30:261–273.

Whiten, A., and R. Ham. 1992. On the nature and evolution of imitation in the animal kingdom : reappraisal of a century of research. *Adv. Study Behav.* 21:239–283.

Associate Editor: M. Doebeli