

Applying evolutionary models to the laboratory study of social learning

Richard McElreath^{a,b,c,*}, Mark Lubell^{b,d}, Peter J. Richerson^{b,c,d},
Timothy M. Waring^{b,d}, William Baum^d, Edward Edsten^c,
Charles Efferson^{b,d}, Brian Paciotti^{b,d}

^a*Department of Anthropology, UC Davis, Davis, CA 95616, USA*

^b*Graduate Group in Ecology, UC Davis, Davis, CA 95616, USA*

^c*Animal Behavior Graduate Group, UC Davis, Davis, CA 95616, USA*

^d*Department of Environmental Science and Policy, UC Davis, Davis, CA 95616, USA*

Initial receipt 17 December 2004; final revision received 6 April 2005

Abstract

Cultural evolution is driven, in part, by the strategies that individuals employ to acquire behavior from others. These strategies themselves are partly products of natural selection, making the study of social learning an inherently Darwinian project. Formal models of the evolution of social learning suggest that reliance on social learning should increase with task difficulty and decrease with the probability of environmental change. These models also make predictions about how individuals integrate information from multiple peers. We present the results of micro-society experiments designed to evaluate these predictions. The first experiment measures baseline individual learning strategy in a two-armed bandit environment with variation in task difficulty and temporal fluctuation in the payoffs of the options. Our second experiment addresses how people in the same environment use minimal social information from a single peer. Our third experiment expands on the second by allowing access to the behavior of several other individuals, permitting frequency-dependent strategies like conformity. In each of these experiments, we vary task difficulty and environmental fluctuation. We present several candidate strategies and compute the expected payoffs to each in our experimental environment. We then fit to the data the different models of the use of social information and identify the best-fitting

* Corresponding author. Tel.: +1 530 752 2660; fax: +1 530 752 8885.

E-mail address: mcelreath@ucdavis.edu (R. McElreath).

model via model comparison techniques. We find substantial evidence of both conformist and nonconformist social learning and compare our results to theoretical expectations.

© 2005 Elsevier Inc. All rights reserved.

Keywords: Social learning; Judgment; Conformity; Cultural evolution

1. Introduction

Unlike most animals, humans acquire large and important parts of their behavioral repertoire via imitation and other forms of social learning. Therefore, students of human behavior seek to understand how individuals acquire beliefs and behavior from their parents, peers, and others. At another level, social scientists attempt to fathom the resulting complex interactions that take place at the level of the society. Whether one is interested in the emergence of political institutions, languages, art, technologies, or moral traditions, these cultural elements all arose over long time periods through the combined effects of many individual-level decisions. Understanding how people use information available from the behavior of others is, therefore, important not only for understanding individual decisions, but also for comprehending patterns of change and variation among human societies. And because the psychological mechanisms that make social learning possible are partly products of natural selection, evolutionary models are necessary to fully understand their design.

In this paper, we use micro-society experiments to investigate the psychological foundations of social learning. Experimental micro-societies (Baum, Richerson, Efferson, & Paciotti, 2004) consist of human participants who repeatedly interact in controlled ways within a laboratory. Over a series of rounds, the participants make decisions that lead to real payoffs, receive feedback, and can access some information about the decisions of their peers. Thus, choices evolve over time, in response to both individual and social learning. Our goal is to test and refine hypotheses, developed through formal models of the evolution of cultural evolution, about how people regulate their reliance on individual and social learning and the structural details of how people use social information. We are interested in (1) how social learning changes in response to the difficulty of a task and (2) how it changes in response to the probability of changes in the environment. We develop computational models for estimating decision-making strategies and the strength of different components of these strategies. We find considerable evidence of social learning, and the details of the strategies used in any particular treatment tend to agree with numerical analyses of strategy efficacy. Nevertheless, individuals sometimes imitate in ways that are not obviously profitable. Our results match some of the model predictions concerning task difficulty and environmental change, but by no means all of them.

The first section of the paper reviews the theory that motivates our experiments and lays out qualitative predictions for how people will respond to changes in different kinds of environmental uncertainty. Then, we present in detail our experimental choice environment and analyze how human participants learn on their own within it. These estimates allow us to

analyze the effectiveness of different social learning strategies within our experimental environment. We present a quantitative analysis of the payoffs to three different social learning strategies and the optimal reliance on each, as functions of two types of environmental uncertainty. Then, we present two experiments that allow access to social information and estimate participant strategies in each. Finally, we relate the findings to the predictions derived from the theoretical literature and our own analysis.

2. The evolution of social learning

In economics and political science, researchers seemingly discovered social learning as a “rational” phenomenon in the early 1990s (Banerjee, 1992; Bikhchandani, Hirschleifer, & Welch, 1992). These models sometimes go by the labels of “herding” or “herd behavior” and, other times, as “informational cascades.” A number of models have been developed that show how ignoring private information and choosing based upon the behavior of others can be optimal. This result is surprising to many social scientists because the intuition dominating the study of judgment and decision making has long been that more objective information improves decisions (Gigerenzer, Todd, & the ABC Group, 1999, demonstrate other ways in which less information use can be optimal).

These models are very similar to models of social learning first derived in biology and anthropology in the 1980s. Formal models by Boyd and Richerson and others derived conditions for natural selection to favor various forms of imitation (Boyd & Richerson, 1985; Rogers, 1988). A rich body of theory now exists arguing both that natural selection will often favor an extensive reliance on imitation and that imitation can lead to unanticipated population-level effects (Henrich & McElreath, 2003, review much of it). A robust result of these models is that social learning never entirely replaces individual learning—no matter how difficult and costly—but that social learning abilities may be broadly adaptive although they severely undermine a society’s ability to track changes in the environment and sometimes lead to the spread of maladaptive behavior (Boyd & Richerson, 1985; Richerson & Boyd, 2004).

A handful of empirical studies have addressed formal models of social learning, whether its general properties (Anderson & Holt, 1997; Kameda & Nakanishi, 2002, 2003) or specific aspects of its design (Apesteguia, Huck, & Oechssler, 2003; Camerer & Ho, 1999; Coultas, 2004; Galef & Whiskin, 2004; Henrich, 2001; Kameda & Nakanishi, 2002; McElreath, 2004; Schotter & Sopher, 2003). Social psychologists, most notably Albert Bandura (Bandura, 1977), of course did a lot of work in 1970s exploring the existence of various cues people use in social learning. However, no formal models of social learning developed in psychology, and most of this work was guided by intuition alone and lead to no enduring analytical work. Economists interested in learning in games have developed a handful of candidate models that may honestly be called “social learning models” (Camerer, 2003). These models are very descriptive and consider only a tiny fraction of the social learning strategies specified in the evolutionary models developed by Boyd and Richerson and others.

We are interested in addressing two variables that recur in many of the formal evolutionary models: (1) the difficulty of learning on one’s own and (2) the frequency of temporal

fluctuations in the payoffs of behavioral options. These can both be thought of as types of environmental uncertainty; however, existing models show that each leads to different qualitative effects on the evolutionarily stable amount of social learning.

First, when it is difficult for individuals to determine the best behavior on their own, a greater reliance on social learning arises at equilibrium (Boyd & Richerson, 1985; Henrich & Boyd, 1998; Rogers, 1988). One way that individual learning might be error prone is if the information available to individuals is of poor quality. For example, if interannual variation in crop yields is large, learners will have difficulty telling if some change in cultivation improves yield in the long run. Crop yield is uncertain, but there is a single best crop. A greater reliance on social learning evolves because social learning can both reduce noise in estimates obtained individually as well as help one avoid costly mistakes that others have already endured.

Second, a principle problem with imitation is that changes in the environment may make past behavior a poor guide to current payoffs. Environments are not perfectly stationary. If the climate or pest populations change, it may no longer be a good idea to plant what one's father planted. Thus, when the frequency of such changes is high, less social learning exists at equilibrium (Boyd & Richerson, 1985; Henrich & Boyd, 1998; Rogers, 1988). Essentially, environmental fluctuation can render useless the adaptive knowledge stored in cultural systems. This fluctuation is another kind of uncertainty, but it reduces rather than increases the amount of social learning at equilibrium.

On the basis of these models, as individual learning becomes more difficult, we expect more social learning, and as the probability of environmental change increases, we expect less social learning. While the evolutionary models do not contain enough psychological detail to say if individuals should facultatively adjust reliance on social learning in different contexts, we think it is reasonable to interpret the intuitions of these models in this way. Based upon cues of difficulty of learning or fluctuation in the environment, people might adaptively regulate their attention to the behavior of others (McElreath, 2004). People may have developed suites of adaptive strategies from which they select, depending on different environmental cues.

Coincidental support for the prediction concerning task difficulty comes from a study of conformity in a perceptual task by Baron, Vandello, and Brunzman (1996), in which conformity appears stronger when the task is made more difficult. Another study that indirectly supports these predictions is a study of consumer choice (Fasolo, 2002), in which participants indicated they would be more likely to use the opinions of their peers in a treatment in which there was no clear best option. McElreath (2004) finds field evidence consistent with qualitative predictions about the difficulty of learning. Galef and Whiskin (2004) study the effects of environmental change on social learning in rats and find results that qualitatively support model predictions. Pingle (1995) constructed a production task in which participants could see the choices of others and finds imitation when changes in the production function (the underlying payoffs to options) were announced. Finally, experiments with groups suggest that conformity is more common in tasks in which there are no obviously correct answers (Kameda, Tindale, & Davis, 2003).

Most of the formal models addressing this problem have modeled imitation of members of a previous generation—parents, elders, older siblings—rather than imitation of peers. In our experiments, there are no naive individuals imitating experienced individuals, as in the

models. Instead, individuals of equal experience have the opportunity to imitate one another. The same predictions hold in this purely horizontal case, however, with some caveats. We demonstrate this in a later section, in which we simulate the performance of different social learning strategies, combined with estimates of how individuals learn individually in our experimental design. First, however, we present the decision-making environment for our experiments and a pure individual learning experiment.

3. Experiment 1: Individual learning

To correctly estimate the use of social learning, we have to take individual learning seriously. In the first experiment, we introduce the task environment and explore patterns of individual learning before introducing the possibility of social learning in later experiments.

3.1. Participants

Thirty-six undergraduates at UC Davis participating in a psychology subject pool took part in this experiment. They participated in groups of 6–10, but each made individual decisions and interacted only with their computer terminal through the course of the experiment. Each participant received course extra credit, in addition to their monetary earnings (see below), for completing the experiment. Experiments lasted about 45 min, and the average earnings in this experiment (as in the other two we present) were US\$6.

3.2. Design

The experiment was programmed using z-Tree (Fischbacher, 2002, the Zürich Toolbox for Readymade Economic Experiments) and administered via computer. All instructions were done via the computer. The protocols and software are available from the authors upon request. The experimental task was framed as a simulated agricultural choice. Each participant faced the decision of planting one of two alternative crops (“wheat” or “potatoes”) each of 20 seasons on each of six sequentially encountered farms (for a total of 120 decisions per participant). On each farm, the mean yield of one crop was higher than the other, but which was higher was random on each farm. Participants were not told the actual means of the crops, as was the case in all the experiments in this article. Each season, participants made a planting decision and was informed of their yield from this decision. Only the most recent yield was ever displayed to the participant, and obviously, no previous yield was displayed in the first season on each farm.

On the first three farms, participants were told that the means were constant on each farm across seasons, but potentially different across farms. In the last three farms, participants were told that which crop was best could change in any given season, and that changes occurred randomly each season, with a chance of 1/20 (communicated as a fraction).

The yields for crop i in each season were drawn from a normal distribution with mean μ_i and variance σ^2 , $y_i \sim N(\mu_i, \sigma^2)$. The variance was the same for both crops, while the mean of

the more profitable crop was 13 units and, that of the less profitable, 10 units. Participants were told at the beginning of the experiment that they would be paid US\$.045 per 10 units of yield, for average total winnings between US\$4 and 8. The stated goal was to maximize their yield by planting the crop with the higher mean yield. We manipulated σ^2 to adjust the difficulty of learning on each farm. When the variance in yield is large, it is harder to learn which of the two crops is best. When the variance is small, the quality of information obtained from planting is much better, and consequently, individual learning more easily discovers the best crop. Each participant planted on farms with three different unknown (but stable) variances in yield: 0.25, 4, and 16. The different variances came in random order for each participant, although the sequence was the same for each participant on the first three and last three farms.

In one extra session using eight participants, we doubled the stakes to check for any large motivation effects. The proportion of correct planting decisions in this session was slightly lower than the other sessions. We concluded that any motivational effect from the size of the stakes was quite minor relative to the variation in behavior in the experiment.

3.3. Results

This decision environment is a variant of the common two-arm bandit with a finite horizon. There is a considerable literature on optimal strategies in such environments (Anderson, 2001; Gittins, 1989); however, it is usually very difficult or impossible to actually compute optimal choices in practice. A smaller number of researchers have investigated how people actually make decisions in these environments (Banks, Olson, & Porter, 1997; Gans, Croson, & Knox, 2003; Horowitz, 1973; Meyer & Shi, 1995), and we know of only one serious study of a Gaussian bandit like our own (Anderson, 2001). There are many possible models (Camerer, 2003), based upon several different views of learning. Our goal here is not to improve upon this literature, but instead to find a robust individual learning model (or models) that we can use as the basis of more complex models in later experiments. The models that we fit in this section are minimally parametric generalizations of some popular candidates (Fig. 1).

To address how participants were using yields to make choices, we fit three different individual learning models to the 4320 planting decisions from this experiment. This allows us to narrow the candidate individual-learning models to use in later fitting exercises. The three models are explained in Table 1. Each of the models uses a different rule to update the estimated mean yields of each crop i in season n , m_i^n , using the yields from each season $y_i^1 \dots y_i^n$. The Bayes 1 model updates the estimate in a Bayesian fashion, assuming that the individual knows the real long-run variance in yield, σ^2 . The Bayes 2 model relaxes this assumption, which means that the sufficient statistic of the estimate of the mean is just the running average of the observed yields (Lehmann, 1983). The final model, Memory Decay, is a generalization of the basic Bayesian model. Instead of the importance of recent information on the estimate being a function of the variance of the estimate and σ^2 , here, it is a parameter γ to be fit from the data itself. When $\gamma=0$, only the most recent information influences the estimate. As γ increases, past observed yields have a greater effect on an individual's estimate of the profitability of a crop.

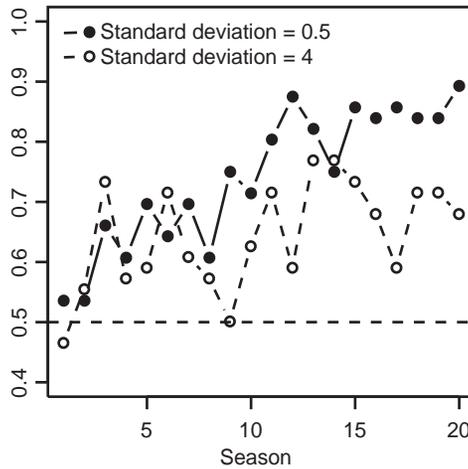


Fig. 1. Effects of standard deviation on task difficulty as shown by the proportion of correct planting decisions in each season, for standard deviations 0.5 (easiest) and 4 (hardest). Standard deviation 2 (not shown) is intermediate between these two trends.

Each model then uses the same functional form, a logit, to model how much each participant cares about differences in the estimated mean yields when choosing a crop to plant in season n . The probability of a participant planting crop i in season n is given by:

$$Pr_n(i) = \frac{\exp(\beta m_i^n)}{\exp(\beta m_i^n) + \exp(\beta m_j^n)}$$

The parameter β captures how much the difference between the estimated means influences choice. When $\beta=0$, choice is random with respect to the estimates of the

Table 1

Models fit to individual learning data and their updating rules for computing the estimate of the mean yield of crop i in season n

Model	Updating rule	Free parameters
Bayes 1	$m_i^n = am_i^{n-1} + (1-a)y_i^{n-1}$ $v_i^n = av_i^{n-1}$ $a = \frac{\sigma^2}{\sigma^2 + v_i^{n-1}}$	β (see main text)
Bayes 2	$m_i^n = \frac{N_i^{n-1}m_i^{n-1} + y_i^{n-1}}{N_i^n}$	β
Memory Decay	$m_i^n = \gamma m_i^{n-1} + (1-\gamma)y_i^{n-1}$	β, γ

The first model, Bayes 1, updates an estimate of the mean m_i and the variance in this estimate, v_i . The second model, Bayes 2, uses the number of samples from crop i , N_i , to compute the running mean. The third model is a parameterized generalization of Bayes 1. In each case, these rules apply when crop i is chosen in round $n-1$. When another crop was chosen in the previous round, each rule is $m_i^n = m_i^{n-1}$.

means. When $\beta = \infty$, the farmer always chooses the crop with the higher mean estimated yield.

The probability model above specifies a likelihood of observing each data point, and we fit each model to the data by finding the values of the parameters that maximize the joint likelihood of observing the data. It is possible to fit these models on an individual-by-individual basis, estimating the strategy that best explains choice for each participant, or across individuals, assuming that each individual is using the same strategy. Using all the data available for each individual, Memory Decay is the best fitting model for 32 of 36 individuals, with an average estimate for γ of 0.11 (maximum=0.50, minimum=0, median=0.065). Bayes 1 is the best fitting model for three individuals, and Bayes 2 for only one individual.

There is too little data for each participant for estimates for each treatment to be reliable, but it is informative to lump together the individuals and fit the models within each treatment. The relative fits of each model may still indicate relative proportions of strategies within the participant pool. Table 2 shows the fits for the three models for the three different variances in yield and the two different fluctuation conditions. The parameter estimates in each case are shown below the Akaike Information Criteria (AIC), Δ value, and Akaike weight (w) of each model. AIC is twice the natural log of the likelihood of observing the data, given the model, plus twice the number of parameters in the model. Thus, smaller AIC values indicate better fits. There is no threshold AIC value that is “good enough.” Fits must be judged relative to one another. The measure Δ is a goodness-of-fit measure analogous to the common R^2 for linear models. For a given model x with minus log-likelihood LL_x , $\Delta_x = 1 - LL_x / LL_{\text{random}}$, where LL_{random} is the fit of a model in which individuals simply guess at each decision (choose randomly). This measure varies from 0, when the fit of model x is the same of the random model, to 1, when the fit of model x is perfect. Therefore, Δ measures the absolute predictive power of a model, compared with a random choice model. Akaike weights (w), in contrast, measure relative fit among the set of considered models. These are computed from the AIC values. The Akaike weight, w_i , for a model i in a set of n models is:

$$w_i = \frac{\exp(-0.5(AIC_i - AIC_{\min}))}{\sum_{j=1}^n \exp(-0.5(AIC_j - AIC_{\min}))},$$

where AIC_{\min} is the smallest AIC value in the set of models considered. Thus, the best-fitting model has the largest w value. One interpretation of Akaike weights is that each indicates the probability that a given model is the correct one. See Burnham and Anderson (2002) for details on these and other measures used to compare models.

The reasons for analyzing data in this way, rather than using common null-hypothesis tests, has been covered many times (e.g., Anderson, Burnham, & Thompson, 2000; Cohen, 1994; Gigerenzer, Krauss, & Vitouch, 2004). Model comparison allows an arbitrary number of competing hypotheses, each of which competes on equal footing. Different specific nonlinear quantitative predictions thus compete to explain the observed data, rather than predictions from only a model that we know a priori to be wrong.

Table 2

AIC, fit relative to a random model (Δ), Akaike weights (w), and parameter estimates for the three individual learning models, by experimental farm standard deviation and probability of fluctuation of means

Standard deviation	0.5		2		4	
Fluctuation	0	0.05	0	0.05	0	0.05
Bayes 1						
AIC	710.84	779.33	781.79	862.04	874.61	932.83
Δ	0.29	0.22	0.22	0.14	0.13	0.07
w	0.14	9.64e-31	3.72e-10	1.56e-40	4.02e-12	2.88e-36
β	0.53	0.63	0.40	0.49	0.28	0.21
Bayes 2						
AIC	710.93	779.37	783.66	863.60	879.52	933.72
Δ	0.29	0.22	0.22	0.14	0.12	0.07
w	0.14	9.43e-31	1.46e-10	7.14e-41	3.44e-13	1.85e-36
β	0.53	0.63	0.39	0.47	0.22	0.17
Memory Decay						
AIC	707.59	641.10	738.37	678.72	822.12	769.16
Δ	0.30	0.36	0.26	0.32	0.18	0.23
w	0.72	1.00e+00	1.00e+00	1.00e+00	1.00e+00	1.00e+00
β	0.53	0.71	0.42	0.54	0.21	0.25
γ	0.05	0.01	0.17	0.05	0.21	0.13

The measure Δ for a model x is defined as $\Delta_x = 1 - LL_x / LL_{\text{random}}$. It gives the relative improvement in fit of the model x , compared with the accuracy of a random model. The weights (w) give the relative fit (one best) of each model, adjusted for number of parameters, to the other models in the analysis.

In every case, Memory Decay is the best fitting model, and only when $\sigma=0.5$ and there is no fluctuation do Bayes 1 and Bayes 2 even approach Memory Decay's fit. The estimates of β show that choice becomes more random with respect to observed payoffs as both variance in yields and the probability of fluctuation increase. The estimates of γ are small in most cases, and below 0.25 in every case. Both decreasing variance and increasing the probability of fluctuation reduce the estimates of γ . This result is quasi-Bayesian: A Bayesian pays less attention to older data when the long-run variance in the data is smaller. New data are informative when the data are not highly variable. Similarly, when the variance in yields here is smaller, γ , the weight given to previous estimates, is smaller. When the means may change each season, previous estimates may become unreliable, and therefore, new data about yield have a stronger influence on the estimate.

3.4. Discussion

Our purpose in the first experiment was to understand individual learning in this decision environment so we can seriously model individual learning in the later experiments that also allow social learning. The Memory Decay model, even accounting for its extra parameter, fits much better than the two Bayesian models. We found that the degree to which individuals are influenced by differences in yield trends downwards (as indicated by lower β estimates) when the variances in yield are high or the environment is not stochastic. We also mapped out the

difficulty of learning the correct crop, as a function of variance in yield. Given the large difference in rate of learning between $\sigma=0.5$ and $\sigma=4$, we chose these two standard deviations as easy and difficult treatments, respectively, for the next two experiments. Given the clear advantage of Memory Decay in predicting choice in these data, it forms the basis of individual learning in the following analyses.

4. Analysis of strategies

In this section, we use simulations to analyze the performance of three alternative social learning strategies in the experimental environment introduced in Experiment 1. This analysis allows us to make specific predictions about which social learning strategies we expect in each experiment to follow, as well as how much we expect participants to use them. The conditions of our analysis exactly mimic those of the experiments we present afterwards.

Unlike the environment in most of the models that we discussed earlier in the paper, our experiments allow only peer-to-peer cultural transmission. The simulations we present here allow us to see how well the predictions about the effects of difficulty of learning and fluctuations in payoffs hold in our modified case. They also allow us to make immediately relevant comparisons of the effectiveness of different social learning strategies.

Many social learning strategies are available to people in natural environments. Our experiments restrict people to strategies that rely upon the frequencies of different alternative behaviors. We outline three different imitation strategies of this type.

Linear imitation. When individuals choose a target individual at random and copy their observed behavior, we refer to this as Linear Imitation. The imitation is linear with respect to each behavior's frequency in the population of potential targets. For example, if two alternative behaviors are present with frequencies 0.6 and 0.4, then linear imitation has a chance 0.6 of copying the first and 0.4 of copying the second. Across iterations of social learning, linear imitation does not change the expected frequencies of behaviors in the population.

We model Linear Imitation in a nested model with the individual learning model fit in the previous section. Let L_i^n be the probability of choosing behavior i in round n from the Memory Decay model. Then, the probability of choosing behavior i in round n when using Linear Imitation is:

$$Pr_n(i) = (1 - \alpha)L_i^n + \alpha \frac{x_i^{n-1}}{N}, \quad (1)$$

where x_i^n is the number of observable target individuals who choose option i in round n and N is the total number of observable targets. The parameter α specifies the strength of reliance on imitation versus individual learning. When $\alpha=0$, the model reduces to the pure Memory Decay model. When $\alpha=1$, the model reduces to pure Linear Imitation.

Confirmation. Another way to use the behavior of a single target individual, without simply copying his behavior, is to practice Confirmation. By Confirmation, we mean keeping

one’s previous behavior, when a randomly chosen target individual also previously chose the same behavior, and relying upon individual judgment otherwise.

We model Confirmation in a nested model, as we do with Linear Imitation. Assuming only two behavioral options, the probability of adopting behavior i in round n , given access to N target individuals, x_i^{n-1} of whom practiced behavior i in round $n-1$, is:

$$Pr_n(i) = (1 - \alpha)L_i^n + \alpha \left(\frac{x_i^{n-1}}{N} \begin{cases} c^{n-1} = i, & 1 \\ c^{n-1} \neq i, & L_i^n \end{cases} + \frac{N - x_i^{n-1}}{N} \begin{cases} c^{n-1} = i, & L_i^n \\ c^{n-1} \neq i, & 0 \end{cases} \right), \quad (2)$$

where c^{n-1} is the individual’s behavior (choice) in round $n-1$.

Conformity. When at least three target individuals are observable, one can do better by using information from each of them. We define Conformity as adopting the majority behavior among a group of targets. When there is no clear majority among the targets, we assume that individuals fall back on individual judgment.

In our two-alternative choice environment, the probability of choosing behavior i in round n is:

$$Pr_n(i) = (1 - \alpha)L_i^n + \alpha \begin{cases} x_i^{n-1} > N/2, & 1 \\ x_i^{n-1} < N/2, & 0 \\ x_i^{n-1} = N/2, & L_i^n \end{cases}. \quad (3)$$

4.1. Comparison of strategies

Which of these strategies is best in our experimental setting, and what is the optimal amount of reliance on each? The theory we mentioned in Section 2 (Henrich & Boyd, 1998, especially) suggests that Conformity is broadly adaptive and likely to perform better than either alternative that we have nominated. However, Confirmation has not yet been analyzed in the thorough way that Conformity has been. Section 2 also suggests that reliance on any social learning strategy should increase with increasing difficulty of the task and decrease with increasing fluctuation in payoffs. To compare the three strategies above, therefore, we conducted simulations to compute the expected payoffs to each strategy, under different values of the variance in yields (σ) and the probability of fluctuation in the means (f), as well as across the range of reliance on each (varying α from 0 to 1). We compute payoffs to groups composed of individuals with a single value if α , because it is easy to show that the payoff function contains a single optimum (Boyd & Richerson, 1985). As long as all individuals have the same ability to learn, there will be a single evolutionarily stable value of this parameter.

The simulations use the exact experimental design described in Experiment 1, except that the virtual participants are in fixed groups of five members and can freely observe the previous choices, but not payoffs, of each other group member in each planting round. We conducted 100,000 simulations at each parameter combination, where each simulation modeled decisions in 20 rounds of planting. Unless noted otherwise, parameters are set at $f=0$ and $\sigma=4$. Performance was measured by the mean payoff over all 100,000 simulations. The values of the individual learning parameters, β and γ , used in the simulation were taken from

the maximum-likelihood estimates from the Memory Decay model in Experiment 1, which are the best guess as to how participants are learning individually. Of course, individual learning may change when social information is introduced. However, as we will demonstrate later, the estimates of these parameters change very little in our subsequent experiments. Simulated participants always access social information in each season (round); however, this will not necessarily be the case in the experiments we present next, where we use both voluntary access and use of social information to estimate strategies.

We found no situation in which Conformity, at its optimal value of α , does not outperform both Linear Imitation and Confirmation. Fig. 2 shows the expected payoffs to Linear Imitation, Confirmation, and Conformity, as functions of the reliance on social learning (the value of the parameter α in the models above). Linear Imitation is never useful, in this environment. The expected payoff to Linear Imitation is always highest when the reliance on social learning is zero. Confirmation and Conformity both lead to gains over both pure individual learning (when $\alpha=0$) and Linear Imitation, but Conformity outperforms both other strategies, provided that the individual uses the optimal value of α . These results are typical of other experimental settings. Conformity leverages the extra information available from multiple target individuals, while neither other strategy does so. Increasing the size of social groups would increase the advantage that Conformity holds.

However, when information from only one target is available, Confirmation is better than Linear Imitation. Simulations with two-person groups, in which only one other individual is observable, confirm that Confirmation's effectiveness generalizes to these smallest possible social groups. Conformity is not possible in these groups, however.

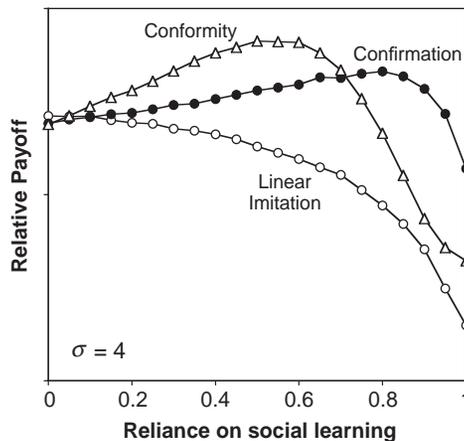


Fig. 2. Simulated relative performance of linear (unbiased) imitation, confirmation, and conformity. Results plotted for simulations of five-person groups, with individual learning based on the empirical estimates. When possible, conformity outperforms confirmation. Both conformity and confirmation outperform linear imitation. Results shown for $\sigma=4$ and $f=0$. The ordering of performance is the same for other values of these parameters. Each point in the graph is the average from 100,000 simulations.

4.2. Optimal amounts of social learning

Confirmation and Conformity both lead to gains in payoff, but the optimal reliance on social learning, measured in the parameter α , will vary as a function of the experimental variables. We demonstrate here how variance in yields and fluctuation in the means lead to the predicted effects that we summarized in Section 2.

Figs. 3 and 4 plot the relative performance of Confirmation and Conformity under different experimental conditions. In both cases, increases in the variance in crop yields make an increased reliance on social learning optimal, while an increase in the probability of fluctuation in the means of the yields makes a decrease in reliance on social learning optimal. These computations verify the relevance of the general predictions from the theory reviewed in Section 2. However, in some cases, this difference is quite small. The results of our next two experiments allow us to address the qualitative nature of these predictions, as well as

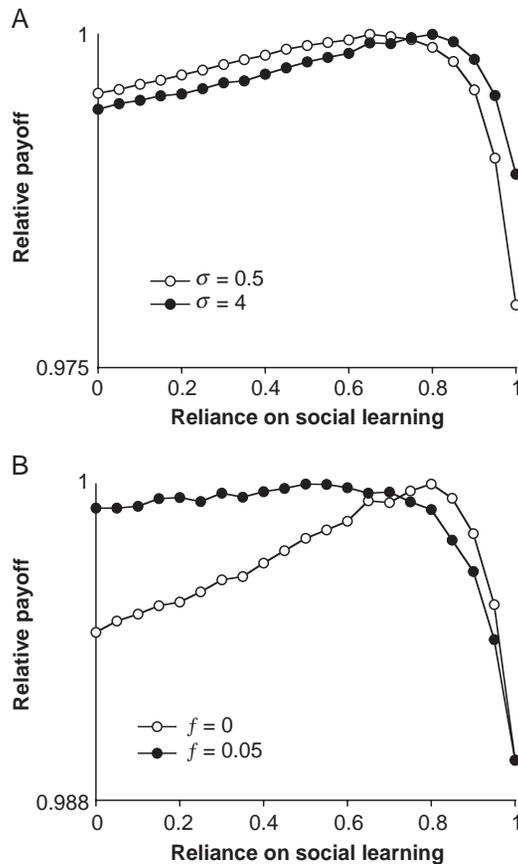


Fig. 3. Simulated relative performance of confirmation under different experimental conditions, for five-person groups. Unless noted otherwise, $f=0$, $\sigma=4$. (A) When the variance in yields increases, more Confirmation-based social learning is optimal. (B) When the environment fluctuates more, less Confirmation-based social learning is optimal.

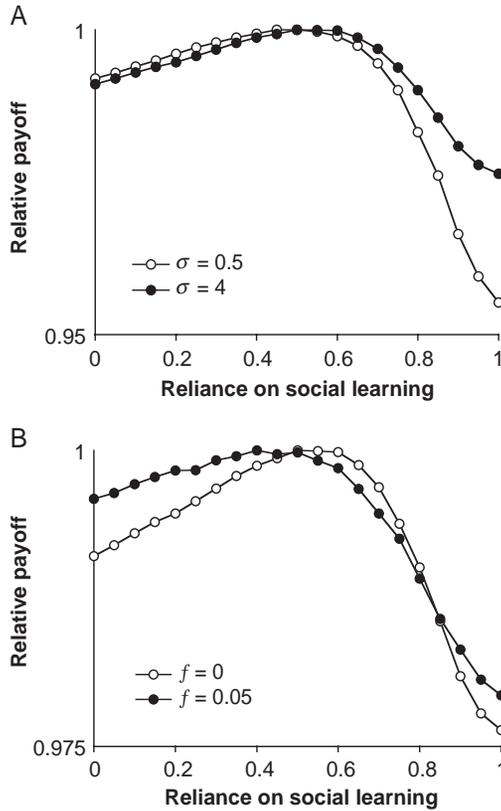


Fig. 4. Relative performance of Conformity-based social learning under different experimental conditions, for five-person groups. Unless noted otherwise, $f=0$, $\sigma=4$. (A) When the variance in yields increases, more Conformity-based social learning is optimal. (B) When the environment fluctuates more, less Conformity is optimal.

measure how calibrated participants are to the decision environment. We do not expect participants to select their strategies optimally, but we do expect detectable shifts in the direction of optimal strategy.

4.3. Predictions for Experiments 2 and 3

Our next two experiments allow access to different amounts of social information. Using the analysis above, we outline a set of predictions for how our estimates of participant strategy will respond to changes in experiment parameters. In Experiment 2, we modify Experiment 1 to allow access to the behavior of a single peer. In Experiment 3, participants have access to the behavior of all group members.

4.3.1. Choice of social learning strategy

We expect participants to rely on Conformity when possible, but to rely on Confirmation in Experiment 2, where information from only one peer is available. Linear Imitation is not

useful in either experiment, and thus, we predict that participants will not use it, provided the cues provided in the experiment lead them to select appropriate strategies.

4.3.2. *Response to variance*

We expect participants to rely on social learning more (as indicated by increased estimates of α) when variance in yield increases, regardless of which strategy they select.

4.3.3. *Response to fluctuation*

We expect participants to rely on social learning less when fluctuation in the mean yields increases, regardless of which strategy they select.

5. Experiment 2: One social target

In the second experiment, we added simple one-model social learning to estimate the participants' tendencies to access this information and how they use it. Recall that our reading of the theory suggests that linear social learning in this purely peer-to-peer laboratory culture is of little value. Instead, we expect to see Confirmation and an increase in reliance upon it when the variance in yield increases and a decrease in reliance upon it when fluctuation in the means is possible.

5.1. *Participants*

Fifty-five undergraduates from a UC Davis psychology subject pool participated in this experiment. None of them had participated in the previous experiment. They participated in sessions of size 6–10, which were divided into anonymous groups of four to six individuals, depending only upon the contingencies of daily attendance. There were 12 groups total: 8 of size 4, 3 of size 5, and 1 of size 6. Participants always knew the actual size of their group, but they never knew the identities of the other people in their group. Participants received course extra credit, in addition to their monetary payments.

5.2. *Design*

This experiment builds upon Experiment 1 by providing one additional piece of information each season, prior to planting. Again, the experiment was programmed using z-Tree (Fischbacher, 2002). All instructions were done via the computer. The protocols and software are available from the authors upon request. The software assigned participants to groups at random and passed decisions among the clients in response to participant behavior. After the first season at each farm, each participant had the option of clicking a button to view the most recent planting decision (but not yield) of one randomly chosen, anonymous member of their own group. Participants were told that members of the same group always experienced the same environment: The means and variances of the crops were the same for all members of a group, at all times, even when the means occasionally switched in the last

three farms. Environmental fluctuations occurred simultaneously within groups, and the participants knew this as well.

5.3. Results

We collected information on the rates at which participants accessed the decisions of other members of their groups (“social information”). We use these data, together with individual planting decisions, to model social learning strategy.

There is impressive variation among participants in the rate that they access social information. Fig. 5(A) plots the distribution of individual click frequencies. Twenty participants never or very rarely accessed social information. The remainder are spread over the entire range of click frequencies. These data alone suggest considerable variation in participant strategy. Fig. 5(B) plots the frequency of clicks for social information averaged across participants but by season (experimental round). The frequency peaks at 0.5 in the second season, the first season social information is available, and declines to just above 0.2 by the final season.

Table 3 shows the estimated changes in odds of accessing social information. These estimates come from a logit probability model with individual fixed effects, allowing us to compare changes across treatments, within participants. Access to social information increases with increasing variance and decreases with increasing probability of fluctuation in the means. A modest interaction term is retained in the optimal AIC model, as well.

While the presentation of the probabilities of access to social information gives hints that social information is of interest to individuals and that participants vary in their interest, they do not tell us much about how participants might be using the information they acquire from other group members. To address this question, we fit the 6360 decisions from this experiment to three candidate models. The first model is the pure individual learning model Memory Decay from the previous experiment. The second model is the Linear Imitation model that we presented in Section 4 (Eq. (1)), which models social learning by introducing one new parameter to the Memory Decay model.

The third model that we fit the data from Experiment 2 is Confirmation (Eq. (2)), also introduced in Section 4, which models participants using social information in a way distinct from copying. For this strategy, we assume the individual checks another participant to see that someone else is doing the same thing as themselves. If the other individual planted the same crop last round as the focal individual, then the focal individual keeps their previous behavior. Otherwise, they rely upon individual learning.

Because participants could see the behavior of only one other individual, Conformity as modeled in Eq. (3) is not possible in this experiment.

Table 4 summarizes the fit and parameter estimates of each of three models, across individuals. Note that these fits are computed assuming that in any round in which a participant does not access social information, the parameter $\alpha=0$, as if they rely exclusively on individual learning for that round. In the easy standard deviation (0.5), the Linear Imitation model fits the best ($w=0.78$) when there is no fluctuation possible, and the Confirmation model is superior when there is fluctuation ($w=0.99$). In the hard

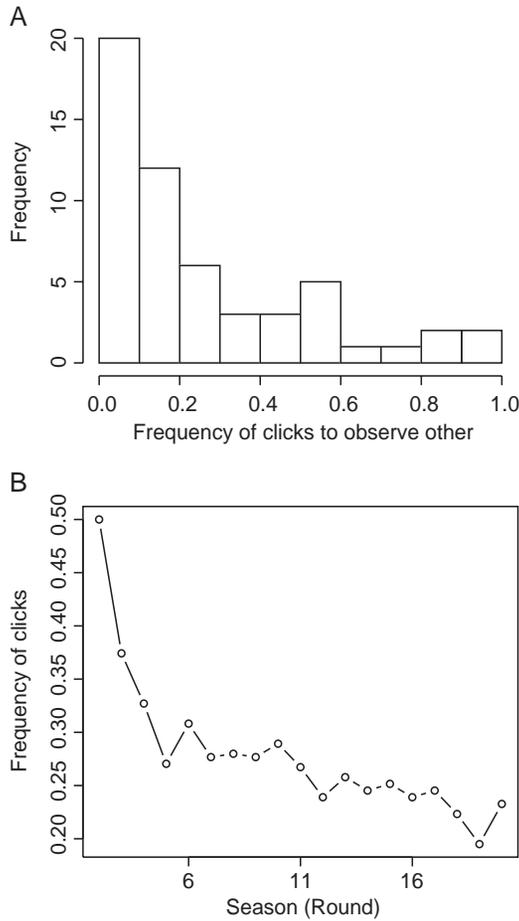


Fig. 5. Access to social information, collapsed across all four experimental farms in each session. (A) Distribution of frequency of clicking to access social information, by individual. (B) Frequencies of clicks to access social information, averaged across individuals, by season. Access peaks in the second season and declines steadily until the final season of each farm. The same group always experienced the same environment: The means and variances of the crops were the same for all members of a group, at all times, even when the means occasionally switched in the last three farms.

standard deviation (4), there is much less dominance of social learning. Each of three models earns good support, although Confirmation consistently does somewhat better than Linear Imitation.

5.4. Discussion

In this experiment, we added only the option of seeing the most recent planting decision of a single anonymous member of one’s own group, who planted under the same conditions. We find that participants choose to view social information slightly more often when the variance in yields is high and when there is no fluctuation in the means through

Table 3

Estimated changes in odds of accessing social information in each experimental round, by standard deviation and probability of environmental fluctuation

Experiment	Factor	Estimate	Standard error	Odds ratio
2	S.D. 0.5	−0.161	0.036	0.852
	Prob fluct 0	0.135	0.035	1.144
	Interaction	−0.087	0.035	0.917
3	S.D. 0.5	−0.087	0.038	0.917
	Prob fluct 0	0.433	0.038	1.541
	Interaction	−0.110	0.038	0.896

These estimates used participant fixed effects to control for individual dispositions and then measure deviations across treatments in a standard two-factor logit probability model. This means that changes within participants are compared across farms, rather than absolute rates of clicking, allowing participants to serve as their own controls. Access to social information increases with increasing variance and decreases with increasing probability of fluctuation in the means. In both Experiments 2 and 3, the full model with both factors and the interaction term provides the best adjusted fit (AIC), compared with simpler models with fewer factors. All observations are included in analysis.

time. We also found substantial evidence of social learning from the model fitting exercise. However, while the social models fit considerably better than the individual learning model for the low-variance farms, this should not be interpreted to mean that all participants were using social information. The social models only differ from the

Table 4

AIC, fit relative to a random model (Δ), Akaike weights (w), and parameter estimates for the three candidate models fit to the data from Experiment 2

Standard deviation	0.5		4	
	0	0.05	0	0.05
Memory Decay				
AIC	1302.06	1353.61	1852.76	1490.37
Δ	0.39	0.41	0.19	0.30
w	0.06	1.97e−05	0.35	0.27
β	0.66	0.79	0.23	0.32
γ	0.01	0.00	0.19	0.00
Linear imitation				
AIC	1296.88	1345.37	1853.30	1491.17
Δ	0.39	0.41	0.19	0.30
w	0.78	1.21e−03	0.27	0.18
β	0.66	0.79	0.24	0.33
γ	0.02	0.00	0.19	0.00
α	0.14	0.18	0.04	0.05
Confirmation				
AIC	1301.02	1331.94	1852.80	1489.04
Δ	0.39	0.42	0.19	0.30
w	0.10	9.99e−01	0.34	0.52
β	0.65	0.76	0.23	0.32
γ	0.00	0.00	0.20	0.00
α	0.17	0.42	0.09	0.17

individual learning model when an individual, in fact, viewed social information. Some individuals rarely did so. By the end of an experimental farm, only about 20% of participants choose to view social information. Thus, the better fit of the social model applies only in those cases, which are overall the minority. When participants did view social information, the evidence indicates that it had a detectable effect on their choices, as indicated by the model fits. Many individuals never or almost never viewed social information, and hence, these individuals must be described as individual learners, despite the better fit of Linear Imitation or Confirmation.

In the General discussion, we interpret the rates of access to social information together with the model fits to evaluate the results of the experiments with respect to the expectations that we developed at the beginning of the paper.

6. Experiment 3: Conformity

While Experiment 2 addresses the most simple kind of social learning possible in our experimental design, rarely do people find themselves in a situation in which they can observe the behavior of only one peer at a time. Experiment 3 was designed to address how participants use social information from more than one individual. Our reading of the formal literature suggests that the use of a majority rule, adopting the most common behavior among models, is more valuable here than simply imitating in the linear or confirmation fashion that was possible in the previous experiment.

6.1. Participants

Forty-nine undergraduates from UC Davis, recruited from classrooms, participated in this experiment. None of them had participated in the previous experiments. They participated in sessions of size 6–10, which were divided into anonymous groups of four to seven individuals, depending only upon the contingencies of daily attendance. There were nine groups total: two of size 4, three of size 5, two of size 6, and two of size 7. Participants always knew the actual size of their group.

6.2. Design

This experiment builds upon Experiment 2 by allowing participants to click a button each season after the first, to view the most recent planting decisions of all other group members. All instructions were done via the computer. The protocols and software are available from the authors upon request. As in Experiment 2, participants viewing social information could not identify individuals by name or number or any other identifying information, nor could they view the payoffs that these individuals received. They simply saw a randomized vector of crop planting choices (wheat, wheat, potatoes, and wheat).

Again, participants were told that members of the same group always experienced the same environment: The means and variances of the crops were the same for all members of a group,

at all times, even when the means occasionally switched in the last three farms. Environmental fluctuations occurred simultaneously within groups, and the participants knew this.

6.3. Results

As in Experiment 2, we collected information about how often participants chose to view the decisions of their peers. Overall, the pattern of clicks is similar to that in Experiment 2: The frequency is highest at the beginning of each farm and stabilizes above zero before season 20. Table 3 summarizes the frequencies of clicks for social information. As in Experiment 2, there is more access to social information in the high-variance farms, especially when fluctuation in the means is absent. There is less access when fluctuation in the means is possible.

While it makes sense to hypothesize that individuals might click more often in larger groups, which contain more information, there is no discernible relationship between group size and frequency of clicks: $n=4$, 0.29; $n=5$, 0.45; $n=6$, 0.36; $n=7$, 0.25.

As with the previous experiments, we analyzed the data from Experiment 3 to determine the ability of different choice models to predict participant planting decisions. The first model that we applied to these data is the Memory Decay model from Experiment 1. This is the baseline individual learning model. The second model that we fit to the data is the Linear Imitation model presented in Section 4 and analyzed already in Experiment 2 (Eq. (1)). The third model that we fit to these data is Confirmation (Eq. (2)), and the fourth, Conformity, introduced in our analysis in Section 3 (Eq. (3)).

We fit these three models to the 5880 decisions made in Experiment 3. Table 5 summarizes the overall model fits. As in previous tables of this kind, we show the overall fit using the AIC, Akaike weights (w), and Δ values. Each model fit is followed by its maximum likelihood parameter estimates. Overall, Linear Imitation and Conformity fit the data much better than either the pure individual learning model or Confirmation. Confirmation does a comparatively poor job everywhere, even compared with pure individual learning. In the absence of environmental fluctuation, Linear Imitation better predicts choice, compared with Conformity, whether variance in yields is high or low. When environmental fluctuation is present, Conformity better predicts choice, when the variance in yields is low ($\sigma=0.5$). When the variance is high ($\sigma=4$), Linear Imitation and Conformity are essentially tied in fit to the data.

Looking at the parameter estimates of α , the reliance on social learning relative to individual learning, the proportions of social learning are overall much higher in this experiment than in Experiment 2. Moving from small variance in yield to large, the estimates of reliance on social learning are relatively much smaller.

Environmental fluctuation—the probability of change in the mean yield of each crop—seems to favor Conformity. In the low-variance farms, Conformity is a much better predictor of choice than is Linear Imitation, provided that the probability of fluctuation is above zero. In the high-variance farms, Conformity ties with Linear Imitation when fluctuation is possible but is far inferior to it when fluctuation is not possible. However, looking at the fits to the individual-by-individual data, when $\sigma=4$ and the probability of fluctuation is .05, Conformity fits the data better than Linear Imitation does. Of 49 participants, the choices of 38 are best

Table 5

AIC, fit relative to a random model (Δ), Akaike weights (w), and parameter estimates for the three candidate models fit to the data from Experiment 3

Standard deviation	0.5		4	
Fluctuation	0	0.05	0	0.05
Memory Decay				
AIC	1753.49	1321.05	1460.98	1518.32
Δ	0.24	0.33	0.18	0.28
w	7.23e–11	1.91e–10	2.54e–03	3.97e–03
β	0.50	0.64	0.21	0.28
γ	0.00	0.03	0.13	0.00
Linear imitation				
AIC	1706.80	1293.83	1449.17	1508.64
Δ	0.26	0.34	0.19	0.29
w	0.99	1.55e–04	0.93	0.50
β	0.49	0.60	0.23	0.31
γ	0.00	0.00	0.11	0.00
α	0.48	0.60	0.23	0.15
Confirmation				
AIC	1755.49	1323.20	1462.96	1519.17
Δ	0.24	0.33	0.18	0.28
w	2.66e–11	6.51e–11	9.42e–04	2.60e–03
β	0.50	0.65	0.21	0.29
γ	0.00	0.06	0.11	0.00
α	0.00	0.00	0.00	0.05
Conformity				
AIC	1716.61	1276.29	1454.40	1508.68
Δ	0.26	0.35	0.18	0.29
w	7.38e–03	1.00	0.07	0.49
β	0.47	0.58	0.21	0.30
γ	0.00	0.00	0.11	0.00
α	0.33	0.54	0.12	0.13

predicted by Conformity, while the choices of 11 are best predicted by Linear Imitation. The evidence suggests that participants are likely to use some strategy approximating Conformity, provided that there is the possibility of fluctuation in the means. Otherwise, there is little evidence of the use of a strategy that integrates the social information in a positive frequency-dependent way.

6.4. Discussion

This experiment allowed participants to see behavior from all other members of their group, and we expected this to lead to conformist crop planting decisions. This expectation was partly upheld, however, not in the absence of environmental fluctuation. When there is no chance of fluctuation in the means of the crops, participants appear to learn socially, but the majority of them are not conformist. We did find, in agreement with the analyses in Section 4, that participants are likely not using a Confirmation strategy.

7. General discussion

The most obvious and least surprising result of our experiments is that many participants used the choices of their peers in making their own choices. It is more surprising that we found evidence of simple Linear Imitation in Experiment 2, where it is little used. Of course, imitation in the one-peer experiment is not much worse than learning on one's own; the payoff difference between imitating a random peer and learning on one's own in this case is not very large. Nevertheless, the sizable proportion (although not the majority) of participants who seem to have used simple imitation deserves an explanation. We imagine two reasons that individuals may imitate when there is no structural feature of the environment or strategy that makes it profitable. First, some individuals learn better than others do. Because some individuals are more likely to arrive at correct behavior faster, for those who imagine themselves slower than average, even Linear Imitation can be profitable. Second, people may be carrying over strategies that are broadly useful in their daily lives into the experiment. It is unreasonable to expect that participants approach experiments as naive yet rational agents. Patterns of imitation behavior in normal life may encourage people to imitate in these experiments, even when there is no apparent advantage to such a strategy.

It is also unclear why participants use Linear Imitation, instead of Conformity, in the no-fluctuation treatments of Experiment 3. On those farms, individuals would have done better had they used Conformity, yet we found little evidence of Conformity there, although we found considerable evidence of it when the environment could fluctuate. Participants are clearly responding to the experimental treatments, and further work will be needed to understand how the cues that the experiment provides activate existing strategies designed for learning in natural settings.

The predictions from theory suggested that individuals would rely more on social learning when (1) the variance in yields was larger and (2) there was no possibility of environmental fluctuation. The clicks to access social information agree with these predictions. Only in Experiment 2, when the standard deviation of yields was small, is there little noticeable effect on the frequency of clicks to access social information (although the measured effect is in the right direction even then). In all other cases, the rates of clicks increase in treatments with higher standard deviations and decrease in those with fluctuation in the means.

However, estimates of reliance on social learning do not generally agree with the predictions. To make the process data (clicks) and model estimates comparable, we need to multiply frequencies of access to social information by the estimated reliances on social learning. We compute the estimated reliance on social learning in each treatment by computing the model-averaged estimate of α in each case, using Akaike weights (w) for weighting the different estimates from different models. The model-averaged estimate of reliance on social learning for a set of n models is:

$$\bar{\alpha} = \sum_i^n w_i \alpha_i^*$$

where α_i^* is the maximum-likelihood estimate of α for model i . For Memory Decay, we set α^* to zero (no reliance on social learning). We compute the total rate of social learning, then,

by multiplying each rate of information access by the model-weighted reliance on that information.

Table 6 summarizes the total estimated frequencies computed in this way. In Experiment 3, the frequencies of social learning decrease, as predicted, when we introduce fluctuation in the mean yields (probability of fluctuation .05). However, in Experiment 2, fluctuation appears to have had the opposite effect: Social learning increases in that case. The effect of increasing the variance in yields is contrary to prediction in every case. In both Experiments 2 and 3, the total estimated frequencies of social learning decrease when the variance in yield increases.

One possible explanation of this counter-theoretical result is that participants are interpreting large variance in the mean yields as environmental unpredictability of the sort introduced by fluctuation. We cannot address this possibility with our data, but new experiments using a different mechanism for manipulating the difficulty of individual learning would help to deal with it. Another possibility is that we simply have the wrong models. All models are simplifications. If the models of social that learning we have considered are missing some structurally important feature of individuals' imitation behavior, then our model estimates will simply tell the wrong story. Notice that the process data, the clicks to access social information, agree with the theory. These process data have the virtue of not being constructed through an intervening model: They are plainly measured. Skeptics of our model estimates may therefore take comfort in the less uncertain process data. Either way, whether because of the details of the difficulty manipulation or the structural inaccuracy if our models, there is a problem to be solved, to reconcile theory with our experimental results.

Many other models are possible. A linear combination of the models that we have presented would explicitly model the idea that individuals sometimes use all of the hypothesized strategies. It is not necessary to include this model in the analysis, however, because the fit relative influences of each strategy would simply be proportional to the Akaike weight of each model, fit separately. Our experiments cannot distinguish between the mixed-strategy and pure-strategy conceptualizations of the results. And because much more data per-individual would be necessary to achieve reliable estimates of individual strategies, we cannot solve the problem in this manner either. Many other functional forms will occur to readers, as

Table 6
Total estimated frequencies of social learning in each experiment and treatment

Experiment	Standard deviation	Probability of fluctuation	Frequency of clicks	Weighted α	Total frequency of social learning
2	0.5	0	0.24	0.13	0.032
	0.5	.05	0.23	0.42	0.099
	4	0	0.32	0.04	0.014
3	4	.05	0.26	0.10	0.025
	0.5	0	0.35	0.48	0.170
	0.5	.05	0.26	0.54	0.140
	4	0	0.42	0.22	0.091
	4	.05	0.29	0.14	0.040

These frequencies come from multiplying the frequency of clicks to view social information by the Akaike-weighted estimated influence of social learning in each treatment.

well. Certainly, the number of specific models of these processes is very large. However, we have kept our family of models small and restricted to those that emerge easily from the general theory. We did this to avoid the temptation to fish for functional forms that fit. In addition, once one begins trying many different functions from different families, it is harder to compare the relative fits of the models. The number of parameters in each model can easily become an inadequate measure of model complexity. There are metrics like minimum description length that are meant to address these concerns; however, their use and interpretation remain controversial.

We found good evidence of individual variation in strategy in all three experiments. Yet, we have made little effort yet to explain this variation. Variation in individual learning ability/skill may explain some of the estimated variation in strategy. It is tempting also to hypothesize about covariance with other individual characteristics. We hope to address the data in this way in later work.

8. Conclusion

By way of conclusion, we offer several cautions. First, these experiments obviously explore only a tiny fraction of the universe of meaningful learning environments and potential strategies available to people. Progress in understanding the design of social learning will come from a body of detailed work fully exploring a number of decision environments and transmission schemes (information structures) while iteratively revising the quantitative models that motivate them. We find work by Tatsuya Kameda and his colleagues (Kameda & Nakanishi, 2002, 2003) inspiring in this regard. It is not enough to simply nominate the existence of a collection of “effects” and test for their existence. Mature predictive models of some depth come about by iteratively building complexity into a research design and the models that it is meant to address.

Second, the depth of this kind of work needs to be balanced by breadth. Replication, of both parameter estimates and general results, both cross-culturally and across cultural domains, is essential. We do not imagine that social learning strategies, which themselves can be learned, are invariant human universals. The strength of conformity, in particular, likely varies cross-culturally and situationally. Students in Western societies are repeatedly admonished to “think for themselves.” It is also important to notice that students, the favorite subjects of psychologists and economists alike, are an odd population to study to understand how people learn. Students in university are trained to learn in particular ways that are unlikely to be representative of most adults. Constructing theories of human nature based on student data is always hazardous, but particularly so in this case.

Even when considering members of a single study population, parameter and strategy estimates from any one sample are notoriously prone to overfitting. Using the estimates that we have developed here to predict the choices of new participants would go a long way to estimating the narrow-sense robustness of our results.

Ultimately, results from laboratory studies like this one need to be validated in naturalistic or quasi-naturalistic settings. Accurate models of individual-level processes can be scaled up to predict large-scale dynamics, much as evolutionary biologists use microevolutionary

models of events in the lives of organisms to understand long-term macroevolutionary trends. Studies like Edgerton (1971); Henrich et al. (2004); Nisbett and Cohen (1996), and Prentice and Miller (1993) remind us of the phenomena that we ultimately intend to understand and provide significant constraints on theory development. More direct applications of microlevel theory to macrolevel problems, like the application of Henrich (2001) of social learning models to data on the spread of technological innovations, demonstrate the relevance of experimental studies to the cultural transformations that we witness in daily life. We think that researchers should not be shy about extrapolating findings in both directions, from the laboratory to the field and visa versa.

Acknowledgments

We thank Joe Henrich, James Fowler, Peter Todd, Gerd Gigerenzer, the members of the ABC research group, and two anonymous reviewers for helpful advice. Thomas Dudey helped with LATEX tricks. This research was funded, in part, by UC Davis junior faculty grants to ML and RM, as well as by National Science Foundation Award Number 0340148 to ML, RM, and PJR.

References

- Anderson, C. M. (2001). Behavioral models of strategies in multi-armed bandit problems. PhD thesis, California Institute of Technology.
- Anderson, D. R., Burnham, K. P., & Thompson, W. L. (2000). Null hypothesis testing: Problems, prevalence, and an alternative. *Journal of Wildlife Management*, *64*, 912–923.
- Anderson, L. R., & Holt, C. A. (1997). Information cascades in the laboratory. *American Economic Review*, *87*, 847–862.
- Apesteguia, J., Huck, S., & Oechssler, J. (2003). *Imitation, theory and experimental evidence*. Available at: <http://www.ideas.repec.org/p/wpa/wuwpex/0309001.html>.
- Bandura, A. (1977). *Social learning theory*. New York: Prentice Hall. Book English.
- Banerjee, A. V. (1992). A simple model of herd behavior. *Quarterly Journal of Economics*, *107* (3), 797–817.
- Banks, J., Olson, M., & Porter, D. (1997). An experimental analysis of the bandit problem. *Economic Theory*, *10*, 55–77.
- Baron, R. S., Vandello, J. A., & Brunsman, B. (1996). The forgotten variable in conformity research: Impact of task importance on social influence. *Journal of Personality and Social Psychology*, *71* (5), 915–927.
- Baum, W. M., Richerson, P. J., Efferson, C. M., & Paciotti, B. M. (2004). Cultural evolution in laboratory microsocieties including traditions of rule giving and rule following. *Evolution and Human Behavior*, *25*, 305–326.
- Bikhchandani, S., Hirschleifer, D., & Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, *100* (5), 992–1026.
- Boyd, R., & Richerson, P. J. (1985). *Culture and the evolutionary process*. Chicago: University of Chicago Press.
- Burnham, K. P., & Anderson, D. (2002). *Model selection and multi-model inference*. New York: Springer-Verlag Telos.
- Camerer, C. F. (2003). *Behavioral game theory: Experiments in strategic interaction*. Princeton: Princeton University Press.
- Camerer, C. F., & Ho, T. (1999). Experience-weighted attraction learning in normal form games. *Econometrica*, *67* (4), 827–874.
- Cohen, J. (1994). The earth is round ($p < .05$). *American Psychologist*, *49*, 997–1000.

- Coultas, J. C. (2004). When in Rome... An evolutionary perspective on conformity. *Group Processes & Intergroup Relations*, 70 (4), 317–331.
- Edgerton, R. B. (1971). *The individual in cultural adaptation: A study of four East African peoples*. Berkeley: University of California Press.
- Fasolo, B. (2002). Multi-attribute decisions online: How people handle conflict among attributes. PhD thesis, University of Colorado at Boulder.
- Fischbacher, U. (2002). *z-Tree: Experimenter's manual*. Zurich, Switzerland: Institute for Empirical Research in Economics.
- Galef, B. G., Jr., & Whiskin, E. E. (2004). Effects of environmental stability and demonstrator age on social learning of food preferences by young Norway rats. *Animal Behavior*, 68, 897–902.
- Gans, N., Croson, R., & Knox, G. (2003). *Customer learning and switching: A two-armed bandit experiment*. Philadelphia, PA: Wharton School, University of Pennsylvania.
- Gigerenzer, G., Krauss, S., & Vitouch, O. (2004). The null ritual: What you always wanted to know about significance testing but were afraid to ask. In D. Kaplan (Ed.), *Sage handbook of quantitative methodology for the social sciences* (pp. 391–408). Thousand Oaks, CA: Sage.
- Gigerenzer, G., Todd, P. M., & the ABC Group (1999). *Simple heuristics that make us smart*. New York: Oxford University Press.
- Gittins, J. C. (1989). *Multi-armed bandit allocation indices*. Oxford: John Wiley & Sons, Inc.
- Henrich, J. (2001). Cultural transmission and the diffusion of innovations: Adoption dynamics indicate that biased cultural transmission is the predominate force in behavioral change. *American Anthropologist*, 103 (4), 992–1013.
- Henrich, J., & Boyd, R. (1998). The evolution of conformist transmission and the emergence of between-group differences. *Evolution and Human Behavior*, 19, 215–242.
- Henrich, J., Boyd, R., Bowles, S., Camerer, C. F., Fehr, E., Gintis, H., & McElreath, R. (2004). Overview and synthesis. In J. Henrich, R. Boyd, S. Bowles, C. Camerer, E. Fehr, & H. Gintis (Eds.), *Foundations of human sociality: Ethnography and experiments in 15 small-scale societies* (pp. 8–54). Oxford: Oxford University Press.
- Henrich, J., & McElreath, R. (2003). The evolution of cultural evolution. *Evolutionary Anthropology*, 12, 123–135.
- Horowitz, A. D. (1973). Experimental study of the two-armed bandit problem. PhD thesis, University of North Carolina, Chapel Hill.
- Kameda, T., & Nakanishi, D. (2002). Cost–benefit analysis of social/cultural learning in a non-stationary uncertain environment: An evolutionary simulation and an experiment with human subjects. *Evolution and Human Behavior*, 23 (5), 373–393.
- Kameda, T., & Nakanishi, D. (2003). Does social/cultural learning increase human adaptability? Rogers' question revisited. *Evolution and Human Behavior*, 24, 242–260.
- Kameda, T., Tindale, R. S., & Davis, J. H. (2003). Cognitions, preferences, and social sharedness: Past, present, and future directions in group decision making. In S. L. Schneider, & J. Shanteau (Eds.), *Emerging perspectives on judgment and decision research*. Cambridge: Cambridge University Press.
- Lehmann, E. L. (1983). *Theory of point estimation*. New York: John Wiley & Sons, Inc.
- McElreath, R. (2004). Social learning and the maintenance of cultural variation: An evolutionary model and data from East Africa. *American Anthropologist*, 106 (2), 308–321.
- Meyer, R. J., & Shi, Y. (1995). Sequential choice under ambiguity: Intuitive solutions to the armed-bandit problem. *Management Science*, 415, 817–834.
- Nisbett, R. E., & Cohen, D. (1996). *Culture of honor: The psychology of violence in the south*. Westview Press.
- Pingle, M. (1995). Imitation vs. rationality: An experimental perspective on decision-making. *Journal of Socio-Economics*, 24, 281–315.
- Prentice, D. A., & Miller, D. T. (1993). Pluralistic ignorance and alcohol use on campus: Some consequences of misperceiving the social norm. *Journal of Personality and Social Psychology*, 64 (2), 243–256.
- Richerson, P. J., & Boyd, R. (2004). *Not by genes alone: How culture transformed human biology*. Chicago: University of Chicago Press.
- Rogers, A. R. (1988). Does biology constrain culture? *American Anthropologist*, 90 (4), 819–831.
- Schotter, A., & Sopher, B. (2003). Social learning and coordination conventions in intergenerational games: An experimental study. *Journal of Political Economy*, 111 (3), 498–529.