Beyond existence and aiming outside the laboratory: estimating frequency-dependent and pay-off-biased social learning strategies

Richard McElreath^{1,*}, Adrian V. Bell², Charles Efferson³, Mark Lubell²,
Peter J. Richerson² and Timothy Waring²

¹*Department of Anthropology, and* ²*Department of Environmental Science and Policy,
University of California Davis, Davis, CA 95616, USA*

³*Institute for Empirical Research in Economics, Blümlisalpstrasse 10, 8006 Zürich, Switzerland*

The existence of social learning has been confirmed in diverse taxa, from apes to guppies. In order to advance our understanding of the consequences of social transmission and evolution of behaviour, however, we require statistical tools that can distinguish among diverse social learning strategies. In this paper, we advance two main ideas. First, social learning is diverse, in the sense that individuals can take advantage of different kinds of information and combine them in different ways. Examining learning strategies for different information conditions illuminates the more detailed design of social learning. We construct and analyse an evolutionary model of diverse social learning heuristics, in order to generate predictions and illustrate the impact of design differences on an organism's fitness. Second, in order to eventually escape the laboratory and apply social learning models to natural behaviour, we require statistical methods that do not depend upon tight experimental control. Therefore, we examine strategic social learning in an experimental setting in which the social information itself is endogenous to the experimental group, as it is in natural settings. We develop statistical models for distinguishing among different strategic uses of social information. The experimental data strongly suggest that most participants employ a hierarchical strategy that uses both average observed pay-offs of options as well as frequency information, the same model predicted by our evolutionary analysis to dominate a wide range of conditions.

Keywords: cultural evolution; social learning; quantitative methods

1. INTRODUCTION

Under a broad definition, social learning is common in nature. The behaviour of conspecifics influences individual behaviour through modification of the environment, emulation of goals and imitation of patterns (cf. Whiten & Ham 1992). This psychological set of distinctions has directed years of research in animal behaviour, especially the study of social learning in non-human apes. Distinguishing between emulation and imitation, and the interaction of the two (Horner & Whiten 2005), has generated a literature testifying to the breadth and diversity of social learning in nature (Fragaszy & Perry 2003).

More recent high-profile experiments with chimpanzees (Whiten *et al.* 2005, 2007) have demonstrated that short-lived traditions can evolve in chimpanzee social groups, backing up studies that claim that behavioural variation among wild populations of chimpanzees are 'cultural' (Boesch & Tomasello 1998; Whiten *et al.* 1999; Boesch 2003). While the finding of short-lived socially transmitted traditions may not be surprising to students of Galef's rat experiments (Galef & Whiskin 1997), the findings suggest that the time may be right to attempt a more

serious exchange between the evolutionary anthropology literature on social learning—which emphasizes a toolbox of social learning strategies, such as majority rule conformity and pay-off-biased learning (Boyd & Richerson 1985; Henrich & McElreath 2003)—and the animal literature—which tends to emphasize the existence or not of culture.

There are at least two good reasons to try. First, non-human animals may also have special-purpose social learning strategies that combine and recombine different kinds of social information, yet usually no effort is made to look for these (Laland 2004). Finding such cases of analogy (or possibly homology, in the case of other apes) would potentiate advances in the general understanding of the evolution of adaptations for processing social information. Second, many biologists and anthropologists remain sceptical of the evidence of animal, and especially great ape, culture (Laland & Janik 2006). This is partly a result of the difficulty of inferring patterns of learning from cross sections of behavioural variation. However, statistical tools developed to study dynamic learning in human groups can be leveraged to study diverse social learning strategies in other animals, as well.

In this paper, we illustrate an approach for analysing different strategies for combining social cues from multiple conspecifics, in less poorly controlled settings.

* Author for correspondence (mcelreath@ucdavis.edu).

One contribution of 11 to a Theme Issue 'Cultural transmission and the evolution of human behaviour'.

We use a stylized evolutionary model to generate broad predictions for which of several candidate strategies we expect to find in nature and under what conditions. We then apply these stylized predictions to a laboratory experiment that allows participants great flexibility in from whom and how they learn. Instead of asking *if* social learning occurs, we develop likelihood models that allow us to ask *how* participants socially learn. While the precise example we present uses very detailed information, the same approach can be applied to more naturalistic contexts, in which incomplete time series or purely cross-sectional data are all that are available.

While neither the appreciation of strategic diversity nor our model-based approach is particularly new in itself, we think the combination is of value. The key insight is that each social learning strategy implies different outcomes, under at least some sets of available information, both for each individual and entire groups of individuals. This is not a new point (Cavalli-Sforza & Feldman 1981; Boyd & Richerson 1985; Galef & Whiskin 1997), but statistical approaches are usually not up to the task of exploring it adequately. Those who do study distinctions among strategies may be inclined to rely upon highly controlled and artificial experiments. Even when an experimenter is clever enough to design a series of treatments that can carefully distinguish among diverse strategies in the laboratory, scientists will still debate the lessons of behaviour in the wild. In order to resolve animal culture debates and gain a more detailed behavioural understanding of social learning, whether in humans or other animals, we will need analytical approaches that do not require precise experimental control of social information. Another reason to develop statistical methods for less controlled contexts is that part of the action in social learning is evolution of behaviour, and experiments that control social information do not allow us to study these population-level effects nor how strategies are adapted to them.

The general approach we suggest is to (i) nominate a series of candidate social learning strategies, (ii) translate each of these into an expression for the conditional probability of behaviour, given an informational context for an individual animal, (iii) use these expressions to generate likelihoods of observing field or laboratory data, and (iv) compare the fits of these strategies to the data with information theoretic criteria, such as Akaike information criterion (AIC) or Bayesian information criterion (BIC). Approaching the problem as a task of discriminating among a toolbox of potential strategies, rather than a task of demonstrating the existence of social learning, may allow all of us to squeeze more from both our experiments and field studies than we previously imagined.

2. MANY WAYS TO LEARN SOCIALLY

There was a time when biology wondered if natural selection occurred. Now no one—within evolutionary biology—seriously questions the existence of natural selection as an evolutionary force. Instead, we debate its relative strength and character in different environmental and biological contexts. Both sexual

(Kokko *et al.* 2006) and social selection (Frank 2006) have generated special literatures of theory and evidence that testify to the subtle diversity of the action of natural selection. One could seriously say that there are many natural selections.

In a similar sense, there are many social learnings. Psychologists and animal behaviourists have long recognized taxonomic distinctions between, for example, social facilitation and imitation (Zajonc 1965). But many of the highest profile publications still address basic existence questions, asking if other animals have human-like social learning and human-like traditions or culture (Whiten *et al.* 2005). These publications are probably taking the right rhetorical approach. Many anthropologists remain unconvinced that chimpanzee or crow culture is much like human culture (Boesch 2003).

However, many scientists have enough interest in the details of social learning in humans, as well as other animals, to step aside the ‘is it human enough?’ debate. As social learning is diverse, it has diverse effects. Some mechanisms generate rather short-lived traditions, if any at all (Galef & Whiskin 1997). Human cultural traditions can be both ephemeral and demonstrate tremendous inertia (Richerson & Boyd 2005), depending in part upon the strategic diversity of social learning and the details of the social context (Cavalli-Sforza & Feldman 1981; Boyd & Richerson 1992). Studying the mechanistic and algorithmic diversity of social learning will be just as important as arguing that it exists, and our hunch is that most researchers in both anthropology and animal behaviour are prepared to move in this direction.

In this section, we briefly review evolutionary work on structurally different social learning strategies. Most of this literature has been concerned with human social and cultural learning (Boyd & Richerson 1985; Henrich & McElreath 2003), but there is no reason these models cannot apply to other organisms (Laland 2004). Before moving on to apply these different strategies to experimental data, we hope to convince the reader that it is worth asking, for example, if chimpanzees also use majority rule social learning or are guided by observed cues of others’ success. While no single strategy is imagined to dominate at all times nor to exist in the absence of individual learning, the dynamic consequences of each strategy can be appreciated most easily by first examining them in isolation.

(a) *Unbiased social learning*

One of the simplest social learning strategies is to select a random target individual and copy his or her behaviour. We call this kind of social learning ‘unbiased’, as it tends to maintain the frequencies of different behaviour (Cavalli-Sforza & Feldman 1981; Boyd & Richerson 1985). One adaptive advantage of unbiased social learning is economizing on individual learning costs (Boyd & Richerson 1985).

(b) *Frequency-dependent social learning*

When individuals can sample more than one conspecific, a large family of frequency-dependent strategies become possible. The most commonly

Figure 1. (a,c) Instantaneous and (b,d) evolutionary dynamics of frequency-dependent and pay-off-biased social learning.

studied of these is *positive frequency dependence*, which preferentially copies the most common behaviour variants in the sample. Such a strategy has very deep intellectual roots, being studied formally at least as far back as 1785, in Condorcet's jury theorem (see Estlund 1994). Evolutionary treatments of positive frequency dependence, 'conformity', emphasize its adaptive value for individuals (Boyd & Richerson 1985; Henrich & Boyd 1998).

Figure 1a,b plots the instantaneous and evolutionary dynamics of positive frequency dependence. In figure 1a, for any frequency of one of two alternative learned behaviours on the horizontal axis, the solid curve gives the expected frequency (or probability of adoption) after social learning. If p is the value on the horizontal axis, then $p + p(1-p)(2p-1)$ is the value on the vertical (Boyd & Richerson 1985—we re-derive this function in §2e). The dashed line illustrates the expected frequency under unbiased social learning. In figure 1b, the evolution of behaviour within a population of learners who practice positive frequency dependence depends on whether the initial frequency of behaviour is below or above one-half. Positive frequency dependence tends to increase the more common variants and decrease the others.

(c) Pay-off-biased social learning

When individuals have information about the pay-offs of others, it is possible to use these cues of success to adaptively bias social learning. Such *pay-off*-, *success*-, or *prestige-biased* social learning can be very individually adaptive, provided cues are reliable, leading to

evolutionary dynamics that can be very similar to natural selection (Boyd & Richerson 1985; Schlag 1998, 1999; Henrich & Gil-White 2001). A key property of these strategies may be their tendency to lead to the copying of neutral or mildly maladaptive behaviour that was initially associated with successful individuals (Boyd & Richerson 1985), but recombination is also a possibility (Boyd & Richerson 2002).

Figure 1c,d plots the instantaneous and evolutionary dynamics of simple pay-off-biased learning. In figure 1c, frequency of trait after social learning as a function of the frequency before social learning is shown. If p is the value on the horizontal axis, then $p + p(1-p)b$ is the value on the vertical axis (derived in McElreath & Boyd 2007, ch. 1). The parameter b determines the strength of pay-off bias and is analogous to a selection coefficient, in genetic evolutionary theory. We plot $b=1/2$ here. The dashed line is again the expectation under unbiased social learning. In figure 1d, the evolutionary dynamics produce a classic logistic growth curve (solid curve). Pay-off-biased social learning tends to increase the frequency of adaptive behaviour, but at the cost of greater information demands.

(d) Integrated social learning

Many mixes of the above kinds of social learning are possible (Laland 2004; Whiten et al. 2004). Aside from the likely possibility that individual asymmetries—age, sex, skill, position in social network—will make some strategies more common among some individuals, strategies can be hierarchically ranked within each individual. Mixes of strategies produce their own

evolutionary trajectories, as well (Henrich 2001). The dashed curve in figure 1*d* is the dynamics of a mix of pay-off bias and positive frequency dependence. For different mixes of these and other strategies, different evolutionary dynamics are expected.

(e) Modelling integrated pay-off-biased and frequency-dependent social learning

While there has been modelling effort devoted to studying linear, unbiased social learning, frequency-dependent social learning and pay-off-biased social learning, to our knowledge no theoretical study has simultaneously examined these options in the same context. Therefore, we finish this section by presenting an extension of existing evolutionary theory that includes frequency-dependent bias, pay-off bias and a hierarchical integration of the two. We construct recursions for the dynamics of genes controlling these different learning strategies, as well as for the frequency of adaptive learned behaviour. We then analyse this gene-culture system in order to understand what environments favour different strategies.

Consider a large population living in a uniform but temporally varying environment. Each individual faces a choice of two discrete behaviours. One of these choices yields a fitness benefit B , a proportion a of the time, yielding an average of aB . The other yields an average $bB < aB$. For each generation, there is a chance u that the better behaviour switches to the other option. These changes cannot be observed by individuals.

Behaviour is acquired via learning, either individually or socially. Individual learning (I) pays an average learning cost in order to determine which option is better. This makes the fitness of an individual learner:

$$W(I) = w_0 + aB - c,$$

where w_0 is baseline fitness from other behaviour and c is the average cost of learning.

Social learning can be unbiased (linear, L), frequency dependent (conformist, C), pay-off biased (S) or pay-off conformity (SC). Linear social learning copies a random adult from the previous generation, resulting in average fitness:

$$\begin{aligned} W(L) &= w_0 + aBq + bB(1-q) \\ &= w_0 + B(aq + b(1-q)). \end{aligned}$$

The frequency of currently optimal behaviour, q , has its own dynamics, which we define below. The important point here is that linear social learning does not transform this proportion in any direct way. On average, it replicates the frequency of optimal behaviour across generations.

Positive frequency dependence, conformity (C), does however transform q . We assume perhaps the simplest conformity heuristic. The learner samples three random adults from the previous generation and then adopts the most common behaviour among these three models. Since the chance that any one model has optimal behaviour is q , the binomial distribution (table 1) allows us to compute the probability of any combination and therefore the probability of the

Table 1. Probabilities of acquiring optimal behaviour via social learning, for the three nonlinear strategies positive frequency dependence (C), pay-off bias (S) and pay-off conformity (SC). The sample column gives all possible samples of three adults. Uppercase letters indicate that the individual sampled received a large (B) pay-off from their behaviour whereas lowercase indicates the opposite. A or a indicates optimal behaviour and B or b indicates non-optimal behaviour. By multiplying each probability of a specific sample occurring by chance in a particular strategy column, one can sum these products to compute an expected probability of acquiring optimal behaviour via a given strategy. In the Pr(sample) column, q is the frequency of optimal behaviour in the entire adult population and a and b are the probabilities of optimal and non-optimal behaviours, respectively, returning large pay-offs.

sample	Pr(sample)	Pr (1 C)	Pr (1 S)	Pr (1 SC)
AAA	$q^3 \cdot a^3$	1	1	1
AAa	$q^3 \cdot 3a^2(1-a)$	1	1	1
Aaa	$q^3 \cdot 3a(1-a)^2$	1	1	1
aaa	$q^3 \cdot (1-a)^3$	1	1	1
AAB	$3q^2(1-q) \cdot a^2b$	1	2/3	1
AAb	$3q^2(1-q) \cdot a^2(1-b)$	1	1	1
AaB	$3q^2(1-q) \cdot 2a(1-a)b$	1	0	0
Aab	$3q^2(1-q) \cdot 2a(1-a)(1-b)$	1	1	1
aaB	$3q^2(1-q) \cdot (1-a)^2b$	1	0	0
aab	$3q^2(1-q) \cdot (1-a)^2(1-b)$	1	2/3	1
ABB	$3q(1-q)^2 \cdot ab^2$	0	1/3	0
ABb	$3q(1-q)^2 \cdot a2b(1-b)$	0	1	1
Abb	$3q(1-q)^2 \cdot a(1-b)^2$	0	1	1
abb	$3q(1-q)^2 \cdot (1-a)(1-b)^2$	0	1/3	0
aBB	$3q(1-q)^2 \cdot (1-a)2b(1-b)$	0	0	0
aBb	$3q(1-q)^2 \cdot (1-a)b^2$	0	0	0
BBB	$(1-q)^3 \cdot b^3$	0	0	0
BBb	$(1-q)^3 \cdot 3b^2(1-b)$	0	0	0
Bbb	$(1-q)^3 \cdot 3b(1-b)^2$	0	0	0
bbb	$(1-q)^3 \cdot (1-b)^3$	0	0	0

conformist learner acquiring optimal behaviour is

$$q_C = q + q(1-q)(2q-1).$$

Using this expression gives us a mean fitness for C,

$$W(C) = w_0 + B(aq_C + b(1-q_C)).$$

Pay-off-biased social learning (S) samples three individuals and adopts the behaviour with the highest average observed pay-off. We compute the expected probability of acquiring optimal behaviour through this heuristic in the same fashion as for conformity: each of the three models sampled has a chance q of having optimal behaviour and each model then has a chance either a or b of displaying a pay-off of B . Thus, the probability of any combination of underlying behaviour and displayed pay-offs can be computed from the binomial distribution (table 1). This results in a chance of acquiring optimal behaviour:

$$\begin{aligned} q_S &= q + q(1-q)(a(2 + b(2 - 3b(1-q) - 4q)) \\ &\quad + a^2(3b - 1)q + b(b(1-q) - 2)). \end{aligned}$$

Figure 2. Sensitivity analysis for the evolutionary model of social learning strategies in the main text. Each row plots the frequencies of five different strategies (individual learning, unbiased social learning, positive frequency dependence, pay-off bias and pay-off conformity) for two-dimensional combinations of parameters. Each individual plot is the frequency of a single strategy after 5000 generations of simulations at all combinations of the two parameters labelled on each axis. (a(i)–(v)) u varied from 0 to 0.5, b varied from 0 to 0.5, while $a = 0.5 + b$. (b(i)–(v)) $a - b$ varied from 0 to 0.5, b again from 0 to 0.5. (c(i)–(v)) u again varied from 0 to 0.5, B from 2 to 10. All other parameters not on axes were fixed at $B/c = 6$, $u = 0.1$, $a = 3/4$, $b = 1/4$, $w_0 = 2$. The most powerful inference from these simulations is that either pay-off bias (S) or pay-off conformity (SC) dominates the population, unless the environment is very unstable and individual learning is too costly, relative to fitness benefits of optimal behaviour.

The fitness of S is therefore,

$$W(S) = w_0 + B(aq_S + b(1 - q_S)).$$

Finally, we consider the integrated strategy pay-off conformity (SC). This strategy attempts pay-off-biased social learning just as S, but falls back on positive frequency dependence whenever observed pay-offs are tied. Just as before, it is possible to compute the expected chance of acquiring optimal behaviour through this heuristic, by using the binomial distribution (table 1). This gives us

$$q_{SC} = q + q(1 - q)(3a(1 - b^2)(1 - q) + 3a^2bq - q(3b - 2) - 1).$$

Again, this implies mean fitness:

$$W(SC) = w_0 + B(aq_{SC} + b(1 - q_{SC})).$$

The dynamics of q are governed by the proportions of each strategy in the population. For proportions $f_I, f_L, f_C, f_S, f_{SC}$, the frequency of optimal behaviour in the next generation in the absence of environmental change is given by

$$q' = f_I + f_Lq + f_Cq_C + f_Sq_S + f_{SC}q_{SC}.$$

Now accounting for environmental change, we arrive at the recursion for the frequency of optimal behaviour in the next generation:

$$q'' = (1 - u_t)q' + u_t(1 - q'),$$

where $u_t \in [0, 1]$ is a random variable indicating whether the environment changed in generation t . This random variable has chance u of being 1, as u is the long-run rate of environmental change.

The complete evolutionary system is very difficult to analyse, because the recursion for q is highly nonlinear. This means there is no guarantee that q even reaches a stationary distribution, and so the fast-slow dynamics approach often employed in these situations (see McElreath & Boyd 2007, ch. 6) is risky. Even if we adopt the fast-slow approach, the implied equilibrium of q is itself the solution to a cubic in q and very difficult to analyse.

Therefore, we adopt a simple simulation approach to analysing this system. We conduct simulations for a large number of parameter combinations in order to map out the conditions that favour different strategies. The fitness expressions and the recursion for q allow us to define a set of difference equations that define the evolutionary dynamics of the system. For any initial frequencies of the strategies and values for w_0, B, c, a, b, u , simulating this system amounts to generating a random variable u_t and recursively computing the frequencies of each strategy after selection. After 5000 simulated generations at each parameter combination, we record the frequency of each strategy. While frequencies could in principle be highly stochastic, fluctuating as selection fluctuates, the results show that taking the final frequency delivers the correct inferences. It also turns out that initial frequencies have no effect on the long-run evolution of the system, allowing us to present simulation results for uniform initial conditions in which all strategies had initially equal frequency.

Figure 2 plots the frequencies of each strategy at simulation end, for two-dimensional sensitivity analyses. Black indicates a frequency of 1, white indicates a frequency of 0 and grey indicates

intermediate frequencies, on a linear gradient. Baseline parameter values in these simulations were $B/c=6$, $u=0.1$, $a=3/4$, $b=1/4$, $w_0=2$. In figure 2*a*, the horizontal axis takes b , the rate of good pay-offs from the non-optimal choice, from 0 to 0.5, holding the value of $a=0.5+b$. Thus, the degree to which the optimal choice is better remains constant, but the absolute level of profitability of both options increase, as one moves left to right on the horizontal axis. The vertical axis takes u , the rate of environmental change, from 0 to 0.6, moving top to bottom.

When u is large, the environment changes rapidly, and individual learning excludes the other strategies (figure 2*a*(i)). When the environment is sufficiently stable, however, either pay-off-biased social learning (S) or pay-off-conformity social learning (SC) excludes the other strategies. When b is small, S excludes SC.

The second row varies the difference between the optimal and non-optimal option, $a-b$, from 0 to 0.5, on the vertical axis. The difference in profitability between the two options interacts only very weakly with the absolute level of profitability, shown again on the horizontal axis. At the extreme limit of $a-b=0$, learning does not pay at all, and so all strategies remain at their initial frequencies (the grey line at the top of the plots in figure 2*b*(ii)–(v)), except for individual learning (I), which is eliminated for trying to learn and paying a direct cost to do so.

The third row of simulations interacts environmental uncertainty, u , with the magnitude of pay-offs, B . The vertical axis is identical to that of the first row, but the horizontal varies B from 2 to 10 (centred on the value $B=6$ that generated the other rows). We can see now that, when B is sufficiently small, individual learning is always excluded, even when the environment is highly unstable. Pay-off-biased social learning, however, excludes the other strategies for these parameter combinations. Pay-off conformity only dominates, as the environment becomes more stable. This stands to reason, as conformity—combined with pay-off bias or not—suffers more from changes in the environment than does pure pay-off bias. To understand this, consider what happens to a conformist just after a change in the environment. Chances are, majority behaviour is suboptimal, and therefore conformity tends to reduce the frequency of optimal behaviour even more. Pay-off bias, however, can still use pay-offs as a cue to optimality.

(f) *Analysis summary*

The most obvious result of this analysis is to emphasize the adaptive significance of pay-off-biased social learning, whether combined with frequency dependence or not. Provided pay-offs can be observed with sufficient accuracy, adopting behavioural options with higher observed average pay-offs excludes other strategies under a wide range of conditions. Unless the environment is extremely stochastic (in which case individual learning dominates) or almost perfectly stable (in which case pure conformity dominates), some kind of pay-off-biased learning is an evolutionarily stable strategy, in our simulations.

The integrated social learning strategy, pay-off conformity, excludes pure pay-off bias when the environment is not too unstable. Being partly

frequency dependent, it needs the optimal behaviour to be the more common behaviour, at least long enough to realize fitness gains. Otherwise, ignoring frequency information is more adaptive. The other factor affecting whether pay-off conformity dominates pure pay-off bias appears to be the magnitudes of a and b , the chances optimal and non-optimal behaviour yield large pay-offs. In the simulations, when $a>1/2$, the integrated pay-off-conformity strategy outperforms pay-off bias alone, holding the difference $a-b$ constant. We are unsure what is causing this advantage. The expression $q_{SC}>q_S$ can be reduced, but it yields a complicated expression that is difficult to interpret. It is also not the whole story, because the average value of q is not described by this condition, and a and b will have large effects on this value.

An interesting feature of pay-off-biased strategies is that they can eliminate individual learning, because any variation among individuals in choice can be used to discriminate good and bad options by pay-offs. All of the nonlinear social learning strategies—positive frequency-dependence, pay-off bias and pay-off conformity—can in fact do this, because their nonlinear effects can, under the right conditions, accomplish the same thing as individual learning.

In §3, we present an experimental design that allows for a large number of different and integrated social learning strategies. In light of these simulations, we expect a heavy reliance on pay-off bias. Also, because the environment is quite stable in the experiment (changing every 15 periods, or a rate of 0.07), the integrated pay-off-conformity strategy should exclude pure pay-off bias. We do not think these exact predictions will describe the results—even simple experiments are much more complex than the theory that motivates them. However, if the theory we have presented here gets at the right economic considerations, then the qualitative results should show a much stronger reliance on pay-off bias than frequency bias.

3. EXPERIMENTAL DESIGN

In order to study the diversity of social learning strategies, we require a decision context complex enough to make both frequency dependence and pay-off bias simultaneously possible. Our social learning experiments create social contexts in which groups of individuals can evolve behavioural traditions, through a combination of their own experience and the available social information. These ‘microsociety’ (Schotter & Sopher 2003; Baum *et al.* 2004) experiments are highly controlled, relative to field studies of social learning, and as a result, we know which social and individual information each participant examines at each time step. Unlike most experiments, however, our experimental groups generate all social information endogenously, without any experimenter deception. This both allows us to examine the emergent properties of social learning and develop statistical methods that can address less controlled natural sources of data.

The experiment allows participants to access both the frequencies of different choices and associated pay-offs, within their own social groups. Over a series of rounds, they may or may not use this information to learn, and we use the complete time series of decisions and records of which

participants access which information in order to test the different models of social learning, pay-off biased or frequency biased.

We have used a similar social decision environment in previous work (McElreath *et al.* 2005), and the environment itself is a social-learning extension of familiar multi-arm bandits used in diverse fields to study individual learning. By using a well-studied decision environment, we can begin with good candidate individual learning models and study the effects of adding different kinds of social information. Our previous experimental studies have omitted pay-off information, and so we could not consider pay-off-biased strategies. And while we have used the statistical approach in our previous papers, we have not previously emphasized the methodological value of the statistics themselves, for analysing data collected in 'wild' contexts.

(a) *Participants*

One hundred and sixty-three participants, students at the University of California at Davis, interacted with one another via a computer network. We recruited participants through an advertisement in the campus newspaper. Participants received between \$5 and \$20 for their participation, based upon their performance. We used no deception in this experiment. Participants read a complete set of instructions and successfully completed a set of test questions about their knowledge of the experiment, before beginning.

(b) *Group structure*

Participants were sorted into random, anonymous groups of four to seven individuals, in sessions of between 8 and 20 participants. Each session was a single experiment on a single date. While participants in the same session made choices in the same room, these participants did not know which of the other participants they were sorted into a group with. Groups were constrained to be always greater than three individuals, in order for frequency bias to be effective, as three neighbours are the required minimum for positive frequency dependence. Depending upon the total number of participants showing up for a given session, group sizes were arranged to create as many groups of four as possible. All remaining participants in that session were placed in a single larger group.

(c) *Decision*

Over a series of 60 periods, 'seasons', each participant made a series of 60 crop choice decisions. These 60 periods were divided into four 'farms' of 15 periods each. These farms served to signal to participants that conditions might have changed. On any given farm, one of two crops, 'wheat' or 'potatoes', had a higher average yield than the other. Across farms, which crop was optimal was determined at random. Thus in each period, each participant chose a single crop to plant and receive a yield from. Yields were summed across all periods, and participants received cash payment so that they earned between \$5 and \$20, depending upon performance. The vast majority of participants earned between \$15 and \$18.

The number of farms and periods in each finesses the trade-offs of (i) having only limited time to keep participants before they grow bored and unmotivated and (ii) desiring the most varied data on learning. Thus the total number of periods, 4×15 , is set by the time constraint. The number of periods per farm is set to maximize information about learning dynamics. If we had a single farm of 60 periods, most of the later periods would add little to nothing to the

analysis, because all participants would be sure of the best option by then, as we have learned from previous experiments (McElreath *et al.* 2005). If a farm is too short, however, we never witness the full dynamics of any learning process. Therefore, guided by pilot experiments and our simulation studies, we decided on 15 periods per farm, as this is the approximate value that maximized our ability to correctly distinguish simulated strategies.

(d) *Social information*

On the first period of each farm, no social information was available. However, on each period after the first, participants could access social information from the most recent period. Participants could examine their own most recent crop choices and resulting yields. Each participant could also examine the most recent crop choices and yields of each member of their own group. This information was displayed on screen in boxes labelled by the type of information. When a participant moused over a box, the information in it was displayed. The experiment software tracked millisecond access to this information, resulting in a time series of information access. This kind of 'mouse-tracking' experiment has been used to great effect in judgement and decision-making research (Payne *et al.* 1993). The order of the rows, yield and crop was randomized for each participant, each period, and the order of neighbours was also randomized. The order of the crop choices at the bottom was also randomized within each participant and period.

(e) *Pay-offs*

Both crops generated pay-offs from normal distributions with the same variance, while the better crop had a mean pay-off of 13 units and the worse 10 units (set from previous experience and simulation study). Participants knew that one crop had a constant higher mean than the other, but had no prior information that would allow them to determine which of the two was better.

The variance of yields was constant within farms but could be either 1/2 or 4, determined randomly but in a way to ensure two farms with a variance of 1/2 and two farms with a variance of 4. The different variances comprise a learning difficulty treatment that we have used in previous experiments (McElreath *et al.* 2005).

(f) *Simulating the experiment*

While there is not space here to describe our simulation in detail, we used the statistical models we will present later to produce simulated experimental play, under a variety of group sizes and other experiment parameters. These simulations simply use the probability models to produce stochastic learning and choice. We then run the data produced through the exact statistical analysis we use on the real data. These simulations allowed us to (i) choose good experimental design parameters and (ii) verify that our statistical analysis works (i.e. recovers true simulated strategies).

4. RESULTS

Like our previous experiments (McElreath *et al.* 2005; Efferson *et al.* 2007), participants learn the optimal crop for each farm, over time. Figure 3 shows the proportion of participants making optimal choices, as a function of period within each farm. The rate of improvement is much faster than in previous experiments, which omitted pay-off information for

Figure 3. (a) Proportion of optimal crop decisions, by round within farm. Vertical lines show 95% profile-likelihood confidence intervals. (b) Proportion of neighbours' crop decisions (circles) and yields (curve) inspected, by round within farm.

neighbours (McElreath *et al.* 2005). The increase between periods 2 and 15 is much smaller than the increase between periods 1 and 2.

Perhaps as a result of the marginal gains in optimality declining after the second period, rates of inspecting the choices (which crop was planted) and yields (how much profit was made last period) of neighbours decline from the second period onward (figure 3). The average rate never falls below a majority of neighbours, however. Note that rates of inspecting yields slightly exceed those for inspecting crop choices. This implies that some participants were using something similar to an elimination by aspects strategy, in which one important cue is used to first narrow down the number of cases one will consider (see Payne *et al.* 1993). In this case, some participants may have first eliminated neighbours to examine crop choices from, by first scanning the yields from the previous period. This would result in the kind of pattern seen in figure 3b. Our statistical analyses in §5 use only the yields and crops actually inspected by each participant, and so take the search strategy as a given. We think the design of the search strategy is a worthwhile question, however. But we doubt such details—truly observing information search—will often be possible in natural settings.

5. ANALYSIS

We adopt a statistical approach that allows us to (i) directly use mathematical models of social learning strategies as statistical models and (ii) evaluate several plausible, non-null statistical models simultaneously. The question is not whether social information is used—few would expect a complete absence of social learning in such a context—but rather how social information is used.

(a) Strategies

We translate each hypothetical learning strategy into an expression that yields the conditional probability of an individual choosing any behavioural option i in any period t , given private information and the social information the individual accessed. Each strategy consists of two parts. The first part is the definition of a recursion for updating the attraction scores of all behavioural options. The second part is a convex

combination of individual choice and the influence of social information.

A large number of meaningfully different strategies can be constructed by varying these two components (Camerer & Ho 1999; Stahl 2000; Camerer 2003). As our purpose in this paper is to illustrate the approach in the simplest manner, we do not explore a large strategy space, but instead restrict ourselves to those nominated by the basic research question and existing evolutionary literature: how do people use frequency-dependent and pay-off-biased social learning, when both are possible?

We examine five different models that combine elements of frequency dependence and/or pay-off bias. First, we define (i) individual learning, (ii) frequency-dependent social learning, (iii) pay-off-biased social learning. We then define hierarchical strategies that combine pay-off-biased learning with the frequency dependence or individual learning: (iv) hierarchical compare means and individual learning and (v) hierarchical compare means and frequency dependence. We do not present analyses of strategies that reverse the hierarchical order of information use, frequency dependence and compare means, for example. These strategies fit very poorly to our data, as will become clear when we examine the fits of each basic model, and so we omit them for simplicity of presentation.

(i) Individual learning

We use a standard, successful reinforcement learning model as the basis of individual updating (Camerer 2003, ch. 6). The attraction score of option i in period $t+1$ is given by

$$A_{i,t+1} = (1 - \phi)A_{i,t} + \phi\pi_{i,t},$$

where ϕ is a parameter determining the weight given to new experience and $\pi_{i,t}$ is the pay-off observed for option i in period t . When option i was not sampled in period t , $\pi_{i,t}=0$. Since there is no reason to expect participants to have strong priors favouring either behavioural option, we set $A_{1,0}=A_{2,0}=0$.

The attraction scores are transformed into probabilistic choice with a 'softmax' choice rule, again typical of the learning in games literature. The probability of choosing option i in period $t+1$ is given by

$$\Pr(i|\mathbf{A}_t, \Theta)_{t+1} = \frac{\exp(\lambda A_{i,t})}{\exp(\lambda A_{1,t}) + \exp(\lambda A_{2,t})},$$

where Θ indicates a vector of all parameters and λ is a parameter that measures the influence of differences between attraction scores on choice. When $\lambda=0$, choice is random with respect to attraction scores. As $\lambda \rightarrow \infty$, choice becomes deterministic, in favour of the option with the higher attraction score.

(ii) Frequency-dependent social learning

To model the family of strategies that use the frequency of behaviour among group members, we modify the learning model above to cue choice by the frequency of options seen. Attractions are updated as before, but choice is given by the rule

$$\Pr(i|\mathbf{A}_t, \Theta)_{t+1} = (1 - \gamma) \frac{\exp(\lambda A_{i,t})}{\exp(\lambda A_{1,t}) + \exp(\lambda A_{2,t})} + \gamma \frac{n_{i,t}^f}{n_{1,t}^f + n_{2,t}^f},$$

where $n_{i,t}$ is the count of neighbours observed to have chosen option i in period t ; γ measures the weight of social information on choice; and f determines how nonlinear frequency dependence is. When $f=1$, imitation is unbiased. When $f>1$, however, more common options have exaggerated chances of being copied, resulting in positive frequency dependence, such as majority rule conformity. When $f<1$, frequency dependence is negative, and more commonly observed options are less likely to be copied.

Since changes in choice feedback to changes in attraction scores, even though this strategy has the same attraction updating recursion as individual learning, reinforcement patterns may be quite different.

(iii) Compare means

This pay-off-biased strategy attends to neighbours' yields and chooses the option with the highest observed mean. It uses the choice rule

$$\Pr(i|\mathbf{A}_t, \Theta)_{t+1} = (1 - \gamma) \frac{\exp(\lambda A_{i,t})}{\exp(\lambda A_{1,t}) + \exp(\lambda A_{2,t})} + \gamma \frac{\bar{\pi}_{i,t}^{100}}{\bar{\pi}_{1,t}^{100} + \bar{\pi}_{2,t}^{100}},$$

where $\bar{\pi}_{i,t} = \sum_j \pi_{i,j,t} / n_{i,t}$ is the mean pay-off observed for option i in period t over all group members j , including oneself. Raising these average pay-offs to a large power creates an approximate step function, so that one or the other option is favoured by the social component of choice. When one or both options are unobserved in period t , this strategy behaves as individual learning. We fix $f=100$ in order to force the model to match our theory, i.e. a threshold behaviour.

(iv) Hierarchical compare means/individual learning

This strategy uses the comparison of choice means and individual updating, but in a manner different from the pure compare means model. Using the distance between estimated means as a cue of uncertainty, the strategy falls back on individual learning (attraction updating) when the means are similar. We use a

Figure 4. The function that determines the reliance on pay-off-biased learning, as a function of the observed difference in means, $\bar{\pi}_{1,t} - \bar{\pi}_{2,t}$. See the description of the hierarchical means/conformity strategy in the text. Solid curve, $\delta=1/10$; dashed curve, $\delta=2$.

symmetrical logistic function to model the change in reliance on pay-offs, as the distance between the observed means increases. Let $Y(\delta, \bar{\pi}_{1,t}, \bar{\pi}_{2,t})$ be the proportion of choice that is driven by individual updating, where δ is a new parameter that determines how quickly reliance on pay-offs decreases, as the difference in observed means increases,

$$Y(\delta, \bar{\pi}_{1,t}, \bar{\pi}_{2,t}) = \frac{2}{1 + \exp(\delta(\bar{\pi}_{1,t} - \bar{\pi}_{2,t})^2)} \equiv Y.$$

Figure 4 plots this function for two values of δ . The probability of choosing i under the hierarchical compare means/individual strategy is

$$\Pr(i|\mathbf{A}_t, \Theta)_{t+1} = (1 - \gamma) \frac{\exp(\lambda A_{i,t})}{\exp(\lambda A_{1,t}) + \exp(\lambda A_{2,t})} + \gamma \left((1 - Y) \frac{\bar{\pi}_{i,t}^{100}}{\bar{\pi}_{1,t}^{100} + \bar{\pi}_{2,t}^{100}} + Y \frac{\exp(\lambda A_{i,t})}{\exp(\lambda A_{1,t}) + \exp(\lambda A_{2,t})} \right).$$

For similar observed means, the individual learning component will dominate the social learning term. Otherwise, the individual will mainly attend to differences in observed means. However, if δ is a very large number, then only a very narrow range of very similar observed means will lead to falling back on individual updating.

(v) Hierarchical compare means/frequency-dependent social learning

This model is like the previous, but falls back on frequency-dependent social learning, as the difference in observed means increases.

$$\Pr(i|\mathbf{A}_t, \Theta)_{t+1} = (1 - \gamma) \frac{\exp(\lambda A_{i,t})}{\exp(\lambda A_{1,t}) + \exp(\lambda A_{2,t})} + \gamma \left((1 - Y) \frac{\bar{\pi}_{i,t}^{100}}{\bar{\pi}_{1,t}^{100} + \bar{\pi}_{2,t}^{100}} + Y \frac{n_{i,t}^f}{n_{1,t}^f + n_{2,t}^f} \right).$$

(b) Fitting strategies to data

The 19 experiment sessions involving 163 participants provided 7900 decisions, under full information conditions that might allow us to distinguish between frequency-dependent and pay-off-biased social learning. We fit the above models to these decisions, producing for each model a negative log likelihood of observing the true data, given the assumption that the model is true: $-\log \mathcal{L}(\mathbf{D}|\mathbf{x}, \boldsymbol{\Theta})$ for a model $\mathbf{x}$ with set of parameters $\boldsymbol{\Theta}$, where $\mathbf{D}$ is the data, a vector of ‘crop’ choices. The likelihood is defined as

$$\mathcal{L}(\mathbf{D}|\mathbf{x}, \boldsymbol{\Theta}) = \prod_t \Pr(D_j|\mathbf{A}_{t-1}, \boldsymbol{\Theta})_t,$$

where $\prod_t$ indicates the product over all rows t . The usual practice in likelihood estimation, and the practice we follow here, is to take natural logs of each conditional probability and then sum these to find $\log \mathcal{L}(\mathbf{D}|\mathbf{x}, \mathbf{P})$:

$$-\log \mathcal{L}(\mathbf{D}|\mathbf{x}, \boldsymbol{\Theta}) = -\sum_t \log \Pr(D_j|\mathbf{A}_{t-1}, \boldsymbol{\Theta})_t.$$

Taking logarithms first results in greater precision, owing to the way most computers handle floating point values. The parameters $\boldsymbol{\Theta}$ are fit via maximum likelihood, and therefore the fitting exercise also yields information on the best estimates of flexible components of the learning rules.

We conducted this fitting exercise, as well as the validating simulations, in R and using the helpful package *bbmle* (Bolker and based on *stats4* by the R Development Core Team 2008; R Development Core Team 2008). All analysis code is available from the corresponding author. We confirmed via simulation that our analysis could recover true parameter values and strategies, when the true strategy was among the set of strategies considered. The validation exercise is helpful, because not all distinct models can be distinguished by all kinds of data (this problem has plagued the individual learning literature, see Camerer 2003, ch. 6).

(c) Comparing models

We compare the fit of the social learning models using Akaike information criteria (Akaike 1974; Burnham & Anderson 2002). Unlike null hypothesis testing, comparing models with Akaike information criteria (AIC—called by Akaike himself simply ‘An information criterion,’ but subsequently renamed by the scientific community), or another information criterion, allows a researcher to assess the relative explanatory power of any number of different competing and plausible models, without favouring any ‘null’ model. AIC is an estimate of the information lost by using any particular model to estimate reality.

The advantages of the information theoretic approach over customary null hypothesis testing has been discussed for several decades (see citations in Cohen 1994; Anderson *et al.* 2000), so we will not repeat them here. Readers should note, however, that there will be no p values in our presentation. Like many statisticians, we do not find much inferential value in p -values, especially when multiple plausible models

are under consideration. AIC and related approaches are becoming increasingly popular in the evolutionary sciences, because they permit more nuanced questions and are not plagued by the same sample size biases of null hypothesis testing (Johnson & Omland 2003). They also allow for more powerful analysis of observational data, collected without precise experimental control.

In order to compare the models, each negative log likelihood from the fitting exercise is transformed into an AIC:

$$\text{AIC}_x = -2 \log \mathcal{L}(\mathbf{D}|\mathbf{x}, \mathbf{P}) + 2k,$$

where k is the number of free parameters in model $\mathbf{x}$. We use the common sample-size-adjusted version of the above, AIC_c (Burnham & Anderson 2002), and this is what we display in our results:

$$\text{AIC}_{c,x} = -2 \log \mathcal{L}(\mathbf{D}|\mathbf{x}, \mathbf{P}) + 2k + \frac{2k(k+1)}{n-k-1},$$

where n is the number of observations to be predicted by the model. The penalty for number of parameters is not arbitrary—it adjusts precisely for the expected overfitting that arises whenever free parameters are added.

AIC can be used to select a single ‘best’ model, if an analyst desires. However, since the ‘true’ model, in all its detail, is certainly not contained in the set of models fit to data, it is perhaps a more productive approach to treat it as a continuous measure of the degree to which each model estimates ‘truth’ (Forster & Sober 1994). AIC estimates the out-of-sample predictive accuracy of each model, and one easy way of ranking these estimates relative to the models in the analysis is by using Akaike weights (Burnham & Anderson 2002). The weight of any model $\mathbf{x}$ is given by

$$w_x = \frac{\exp(-\frac{1}{2} \Delta_x)}{\sum_j \exp(-\frac{1}{2} \Delta_j)},$$

where $\Delta_x = \text{AIC}_x - \text{AIC}_{\min}$, the difference between the AIC of model $\mathbf{x}$ and the smallest AIC in the set of compared models. For the best-fitting model with the smallest AIC, $\Delta = 0$. These weights are numbers between 0 and 1 that estimate the relative likelihoods of each model being the best model in the set.

A useful way we have found to explain this approach is to consider a horse race. There are many horses in each race, and while the fastest horse will not always win, it usually will. If the best horse loses, it should not usually lose by much. Thus both the rank of finishes—which horse was first, second, etc.—and the time differences in finishes are informative. In the same way, the true model may not always fit the data best (just as ‘significant’ p values do not always identify important effects). But it will usually have a high Akaike weight, even if not the highest. So, just as a photo finish tells you that it is difficult to say, without another race, which of two horses is faster, when two models have very similar Akaike weights, there is uncertainty as to which would make the best out-of-sample predictions. When one model has an Akaike weight much larger than the others, however, we can be confident that it is the best of the models considered.

Table 2. Comparison of social learning models fit to experimental data. AIC_c indicates the adjusted fit of each model to the entire sequence of choices, for each subject, after accounting for model complexity (number of parameters). Models are ordered from best fit to worst. The Akaike weights estimate the proportion of evidence in favour of each model. (Meanings of parameters: $\phi \in [0, 1]$, strength of attraction updating; $\lambda > 0$, influence of attraction differences of choice; $\gamma \in [0, 1]$, weight given to social information; $\delta > 0$, decline in probability of pay-off-biased imitation, as difference in observed means increases. HCMFD, Hierarchical compare means/frequency dependence; HCMINDIV, Hierarchical compare means/individual learning.)

model	AIC_c	Akaike weight	parameter estimates				
			ϕ	λ	γ	δ	f
HCMFD	6519.59	≈ 1	0.6605	0.1645	0.3365	3.210	1.953
HCMINDIV	6918.89	< 0.001	0.4620	0.1917	0.1400	4.982	n.a.
compare means	6924.23	< 0.001	0.4611	0.1921	0.1239	n.a.	n.a.
frequency dependent	6929.34	< 0.001	0.4998	0.1814	0.1349	n.a.	3.396
individual	7004.25	< 0.001	0.4382	0.1866	n.a.	n.a.	n.a.

Table 2 presents the AIC_c , Akaike weight and parameter estimates for each model, sorted from best to worst fitting model. The bulk of evidence favours model 5, hierarchical compare means/frequency dependence. While there is no doubt about heterogeneity among participants, the strength of this result leaves little doubt that any of the simpler strategies accounts for any sizeable fraction of participants. The maximum-likelihood estimate for f , the degree of positive or negative frequency dependence, is just under 2, indicating mild positive frequency dependence or conformity (figure 5). The maximum-likelihood estimate of δ (not shown in figure) produces a steep fall-off in reliance on pay-off bias for a distance above approximately 1 unit. We caution that there is uncertainty in these estimates, but emphasize that a model with δ fixed to a large value, say 100, does not produce a better fit, even accounting for the reduction of one parameter.

Many readers may wonder what proportion of variance in choice is explained by the best-fitting model. As is usual with binomial models, there is no true equivalent of R^2 , the proportion of variation explained by the fit model. However, it is possible to construct an analogue that compares the raw likelihoods of each model to a random choice model. A random choice model just chooses randomly at each time t . Over 7900 choices, this model will always have a negative log likelihood of $-7900 \times \log(1/2) = 5475.863$. This is a reasonable benchmark for the worst any model can do, predicting the data. The negative log likelihood of the best-fitting model is 3259.792. Therefore, an analogous calculation of the variance explained by any model x is $1 - \log(\mathcal{L}(\mathbf{D}|x))/\log(\mathcal{L}(\mathbf{D}|\text{random}))$. In our case, $1 - 3259.792/5475.863 = 0.4047$. For the second-best model, $1 - 3459.442/5475.863 = 0.3682$. These measures do not account for model complexity, but they do provide a rough guide to additional raw variance explained by the best model. We caution, however, that substantial components of choice may be truly random, and therefore any behavioural model will fail to achieve a negative log likelihood of zero. In cases in which measurement error is possible, as in field studies or data coded from video, measurement error will also make it impossible for even the true model to achieve a negative log likelihood near zero.

Figure 5. Maximum-likelihood estimate of strength of positive frequency dependence for the best-fitting model, hierarchical compare means/frequency dependence. Solid curve indicates estimated probability of copying a choice, given its frequency in the group. Dashed line indicates same probability under $f=1$, unbiased social learning.

6. DISCUSSION

We have analysed an interdependent time series of profit-oriented choice behaviour in humans. Our experiment did not precisely control the social information available to each participant. Instead, we allowed all social information to arise endogenously, through the actual behaviour and information seeking of participants. While one major tradition in laboratory experiments frowns upon such a design, we consider it an asset, for two reasons.

First, if the study of social learning is ever to link the psychological to the population level, statistical techniques that can accommodate observational and noisy data are needed. The model comparison approach we adopt in this paper is general to any set of strategies a researcher might imagine. Caution is needed to ensure that the kind of data available can discriminate among the possible models. But provided the different models are identifiable in this way, the likelihood-based information criteria can quantify the relative explanatory power of different hypotheses. These dynamic models can then be reasonably asked to produce out-of-system predictions that provide

another avenue of disconfirmation. By contrast, effect sizes from ANOVA cannot reasonably be expected to predict out-of-experiment effects, because no genuine model of learning is present.

Second, the emergent population-level consequences of social learning can only be studied where the experimenter allows them to occur—in settings in which social information itself is not controlled experimentally. This advantage is twofold. Being able to study population-level effects, such as the emergence of traditions or rates of diffusion, is important. But in a cultural species, such as humans and possibly other species, social learning strategies themselves are probably adapted to a cultural environment (Henrich & McElreath 2007). Thus, it will eventually be difficult to study the functional design of strategic social learning without appreciating the cultural environment it is adapted to. This will be true even (especially) if learning strategies themselves are culturally transmitted, because the population will exert downward causation on individuals' learning strategies.

The major scientific finding of our analysis of the experiment is that our human participants relied heavily on pay-off-biased social learning, as predicted by the evolutionary model. We think predictions generated by an economic, rather than evolutionary model, would make similar predictions, provided social information was endogenous to the model. When there is no additional cost to access pay-offs and the information is subject to no error, as in this experiment, then it is no surprise perhaps that a successful strategy will attend to pay-offs. What might be more counter-intuitive is the hierarchical combination of pay-off and frequency biases. The evidence strongly suggests that our participants used a strategy akin to: (i) Are the two choices' pay-offs similar on average? (ii) If yes, which is more common? (iii) If no, which has the higher average pay-off?

It also worth noting that participants did not require any training time to learn to attend strongly to pay-offs—they did so from the first period when social information was present. We make no strong claims about the source of these strategies. Social learning strategies may of course themselves be learned socially, and we have wondered about the effects of this in previous experiments (McElreath *et al.* 2005). Indeed, there is likely hidden strategic variation among participants. Our analysis approach, fitting a single model at a time to the entire set of data, is a common approach, because rarely do we have enough data on each participant to reliably distinguish differences in strategy. However, in principle, the statistical methods here do not require one to conduct the analysis this way. Each participant can be analysed separately, or a series of fixed effects parameters can be used to statistically model individual differences. In the analysis here, the overwhelming support of the best model implies little strategic variation that could be detected by the considered models. However, we do not think this means all participants used the same strategy, merely that we have not modelled the kind of differences that exist.

A common reaction, both by ourselves and our colleagues, to experiments with students is to be sceptical of the generality of the results. True, university students are a special population that is likely not typical of the human species. However, no single population will likely be representative of the human species. That is, every culture and subculture may be a special case. We think there are serious limits to how much we can generalize from experiments with students. But we also think that being able to explain learning in any case is an advance. Just as studying the evolution of beetle larva in the laboratory does not tell us exactly how evolution works in any other species (or even in wild beetles), the clarity of the results does generate insights than can transfer across cases. Our feeling is that no one should conclude the human species is just like university students, anymore than one should conclude all insects have the evolutionary dynamics of flour beetles. But nor should one ignore flour beetles, as if their evolution is not worth explaining. University students are real people with real learning strategies, and being able to model this learning is worthwhile.

It is always possible that another, unconsidered, strategy is a better description of the social learning process. The same weakness is common to all analytical approaches, however, and we caution readers not to consider this a flaw special to the information criterion and model comparison approach. But despite the strong weight of evidence for this strategy here, we think there is no substitute for replication and the variation of experimental design in order to test the robustness of a result. Both our experiments and theoretical analysis are special, like all experiments and models. Whatever the source of social learning strategies—cultural or genetic or (likely) both—the strategies we find in our experiments certainly did not evolve in the laboratory. And however useful simple evolutionary models are for exploring the logic of population dynamics, they cannot and do not attempt to replicate reality. We have emphasized the generality of the statistical approach, as it is not tied to any particular experiment or set of predictions, but it is worth noting key assumptions of both the experiments and models.

First, the experimental environment we have used provides highly accurate (noise free) pay-off information, whereas real social environments certainly do not. In addition, real social environments may provide cues of success, but these cues will often be integrations of the contributions of many separate behaviours. For example, if someone in your town is healthy, is it a result of her diet, her religion or her close bonds with kin? This integrated nature of cues of success means that people may copy many traits from successful or prestigious individuals, with potentially important effects on cultural dynamics (Boyd & Richerson 1985; Henrich & Gil-White 2001). Relevant to our experimental results and the prediction of the model that pay-off bias would dominate, it may be that the clear advantage of pay-off bias depends upon the ability to know that any cue of success arises from a particular behaviour. If not, other forms of social learning may be

more competitive. Some of our ongoing experiments explore this consideration.

Second, our evolutionary analysis is built upon a number of existing models of the evolution of social learning (Boyd & Richerson 1988, 1995, 1996; Rogers 1988; McElreath & Strimling 2008). By doing so, it is comparable to these models, but also considers a fairly special life history. In all of these models, generations are barely overlapping: adults survive only long enough to be imitated. Grandparents never survive to be imitated. There is no population structure, including within the biological family, and therefore any effects of gene-culture covariation are ignored (see however McElreath & Strimling 2008). While this kind of model provides perhaps the purest evaluation of the logic and economics of social learning strategies, actual strategies may have evolved (culturally or genetically) under rather special conditions or in order to exploit overlapping generations. If so, the inferences derived from these models will be misleading. How they will be misleading is hard to say, until more social learning theory exploring population structure and overlapping generations appears.

This lacuna of theory aside, the existing evolutionary literature is sufficient to motivate the search for positive frequency dependence and kinds of pay-off bias in other apes, if not crows, whales and rats. In the search for the psychological differences that make human cultural evolution qualitatively different from that of other animals, the existence of frequency-dependent and refined pay-off bias is often ignored. For example, experiments in which apes see three ape demonstrators access food through a two-action problem, with two demonstrators performing one action and the third another, will produce data that can estimate the magnitude of positive frequency dependence.

This research was funded by the National Science Foundation.

REFERENCES

- Akaike, H. 1974 A new look at the statistical model identification. *IEEE Trans. Autom. Control* **19**, 716–723. (doi:10.1109/TAC.1974.1100705)
- Anderson, D. R., Burnham, K. P. & Thompson, W. L. 2000 Null hypothesis testing: problems, prevalence, and an alternative. *J. Wildl. Manage.* **64**, 912–923. (doi:10.2307/3803199)
- 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. *Evol. Hum. Behav.* **25**, 305–326. (doi:10.1016/j.evolhumbehav.2004.05.003)
- Boesch, C. 2003 Is culture a golden barrier between human and chimpanzee? *Evol. Anthropol.* **12**, 82–91. (doi:10.1002/evan.10106)
- Boesch, C. & Tomasello, M. 1998 Chimpanzee and human culture. *Curr. Anthropol.* **39**, 591–604. (doi:10.1086/204785)
- Bolker, B. and based on stats4 by the R Development Core Team 2008 *bbmle: tools for general maximum likelihood estimation*. R package v. 0.8.5.
- Boyd, R. & Richerson, P. J. 1985 *Culture and the evolutionary process*. Chicago, IL: University of Chicago Press.
- Boyd, R. & Richerson, P. 1988 An evolutionary model of social learning: The effects of spatial and temporal variation. In *Social learning: a psychological and biological approach* (eds T. Zentall & B. G. Galef), pp. 29–48. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Boyd, R. & Richerson, P. J. 1992 How microevolutionary processes give rise to history. In *Evolution and history* (ed. M. Niteki), pp. 149–178. Chicago, IL: University of Chicago Press.
- Boyd, R. & Richerson, P. 1995 Why does culture increase human adaptability? *Ethol. Sociobiol.* **16**, 125–143. (doi:10.1016/0162-3095(94)00073-G)
- Boyd, R. & Richerson, P. J. 1996 Why culture is common, but cultural evolution is rare. In *Evolution of social behaviour patterns in primates and man* (eds W. G. Runciman & R. I. M. D. John Maynard Smith) *Proc. British Academy*, pp. 77–93. Oxford, UK: Oxford University Press.
- Boyd, R. & Richerson, P. J. 2002 Group beneficial norms spread rapidly in a structured population. *J. Theor. Biol.* **215**, 287–296. (doi:10.1006/jtbi.2001.2515)
- Burnham, K. & Anderson, D. 2002 *Model selection and multi-model inference: a practical information-theoretic approach*. Berlin, Germany: Springer.
- Camerer, C. 2003 *Behavioral game theory: experiments in strategic interaction. The Roundtable Series in Behavioral Economics*. Princeton, NJ: Princeton University Press.
- Camerer, C. F. & Ho, T. 1999 Experience-weighted attraction learning in normal form games. *Econometrica* **67**, 827–874. (doi:10.1111/1468-0262.00054)
- Cavalli-Sforza, L. L. & Feldman, M. W. 1981 *Cultural transmission and evolution: a quantitative approach. Monographs in Population Biology*, vol. 16. Princeton, NJ: Princeton University Press. [L. L. Cavalli-Sforza and M. W. Feldman. ill.; 23 cm. Includes index.]
- Cohen, J. 1994 The earth is round ($p < 0.05$). *Am. Psychol.* **49**, 997–1000. (doi:10.1037/0003-066X.49.12.997)
- Efferson, C., Richerson, P. J., McElreath, R., Lubell, M., Edsten, E., Waring, T. M., Paciotti, B. & Baum, W. 2007 Learning, productivity, and noise: an experimental study of cultural transmission on the bolivian altiplano. *Evol. Hum. Behav.* **28**, 11–17. (doi:10.1016/j.evolhumbehav.2006.05.005)
- Estlund, D. M. 1994 Opinion leaders, independence, and condorcet's jury theorem. *Theory Decision* **36**, 131–162. [In English.] (doi:10.1007/BF01079210)
- Forster, M. R. & Sober, E. 1994 How to tell when simpler, more unified, or less *ad hoc* theories will provide more accurate predictions. *Br. J. Phil. Sci.* **45**, 35.
- Fragaszy, D. & Perry, S. (eds) 2003 *The biology of traditions: models and evidence*. Cambridge, UK: Cambridge University Press.
- Frank, S. A. 2006 Social selection. In *Evolutionary genetics: concepts and case studies* (eds C. W. Fox & J. B. Wolf), pp. 350–363. Oxford, UK: Oxford University Press.
- Galef, B. G. & Whiskin, E. 1997 Effects of social and asocial learning on longevity of food-preference traditions. *Anim. Behav.* **53**, 1313–1322. (doi:10.1006/anbe.1996.0366)
- Henrich, J. 2001 Cultural transmission and the diffusion of innovations: adoption dynamics indicate that biased cultural transmission is the predominate force in behavioral change. *Am. Anthropol.* **103**, 992–1013. (doi:10.1525/aa.2001.103.4.992)
- Henrich, J. & Boyd, R. 1998 The evolution of conformist transmission and between-group differences. *Evol. Hum. Behav.* **19**, 215–242. (doi:10.1016/S1090-5138(98)00018-X)
- Henrich, J. & Gil-White, F. 2001 The evolution of prestige: freely conferred status as a mechanism for enhancing the benefits of cultural transmission. *Evol. Hum. Behav.* **22**, 1–32. (doi:10.1016/S1090-5138(00)00071-4)

- Henrich, J. & McElreath, R. 2003 The evolution of cultural evolution. *Evol. Anthropol.* **12**, 123–135. (doi:10.1002/evan.10110)
- Henrich, J. & McElreath, R. 2007 Dual inheritance theory: the evolution of human cultural capacities and cultural evolution. In *Oxford handbook of evolutionary psychology* (eds R. Dunbar & L. Barrett), pp. 555–570. Oxford, UK: Oxford University Press.
- Horner, V. & Whiten, A. 2005 Causal knowledge and imitation/emulation switching in chimpanzees (*Pan troglodytes*) and children (*Homo sapiens*). *Anim. Cognit.* **8**, 164–181. (doi:10.1007/s10071-004-0239-6)
- Johnson, J. B. & Omland, K. S. 2003 Model selection in ecology and evolution. *Trends Ecol. Evol.* **19**, 101–108. (doi:10.1016/j.tree.2003.10.013)
- Kokko, H., Jennions, M. & Brooks, R. 2006 Unifying and testing models of sexual selection. *Annu. Rev. Ecol. Evol. Syst.* **37**, 43–66. (doi:10.1146/annurev.ecolsys.37.091305.110259)
- Laland, K. N. 2004 Social learning strategies. *Learn. Behav.* **32**, 4–14.
- Laland, K. N. & Janik, V. M. 2006 The animal cultures debate. *Trends Ecol. Evol.* **21**, 544–547. (doi:10.1016/j.tree.2006.06.005)
- McElreath, R. & Boyd, R. 2007 *Mathematical models of social evolution: a guide for the perplexed*. Chicago, IL: University of Chicago Press.
- McElreath, R. & Strimling, P. 2008 When natural selection favors imitation of parents. *Curr. Anthropol.* **49**, 307–316. (doi:10.1086/524364)
- McElreath, R., Lubell, M., Richerson, P. J., Waring, T. M., Baum, W., Edsten, E., Efferson, C. & Paciotti, B. 2005 Applying evolutionary models to the laboratory study of social learning. *Evol. Hum. Behav.* **26**, 483–508. (doi:10.1016/j.evolhumbehav.2005.04.003)
- Payne, J. W., Bettman, J. R. & Johnson, E. J. 1993 *The adaptive decision maker*. Cambridge, UK: Cambridge University Press.
- R Development Core Team 2008 *R: a language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Richerson, P. J. & Boyd, R. 2005 *Not by genes alone: how culture transformed human biology*. Chicago, IL: University of Chicago Press.
- Rogers, A. R. 1988 Does biology constrain culture? *Am. Anthropol.* **90**, 819–831. (doi:10.1525/aa.1988.90.4.02a00030)
- Schlag, K. H. 1998 Why imitate, and if so, how? *J. Econ. Theory* **78**, 130–156. (doi:10.1006/jeth.1997.2347)
- Schlag, K. H. 1999 Which one should I imitate? *J. Math. Econ.* **31**, 493–522. (doi:10.1016/S0304-4068(97)00068-2)
- Schotter, A. & Sopher, B. 2003 Social learning and coordination conventions in intergenerational games: an experimental study. *J. Polit. Econ.* **111**, 498–529. (doi:10.1086/374187)
- Stahl, D. O. 2000 Rule learning in symmetric normal-form games: theory and evidence. *Games Econ. Behav.* **32**, 105–138. (doi:10.1006/game.1999.0754)
- Whiten, A. & Ham, R. 1992 On the nature and evolution of imitation in the animal kingdom: reappraisal of a century of research. In *On the nature and evolution of imitation in the animal kingdom: reappraisal of a century of research* (eds P. Slater, J. Rosenblatt, C. Beer & M. Milinski), pp. 239–283. New York, NY: Academic.
- Whiten, A., Goodall, J., McGrew, W., Nishida, T., Reynolds, V., Sugiyama, Y., Tutin, C., Wrangham, R. & Boesch, C. 1999 Cultures in chimpanzees. *Nature* **399**, 682–685. (doi:10.1038/21415)
- Whiten, A., Horner, V., Litchfield, C. & Marshall-Pescini, S. 2004 How do apes ape? *Learn. Behav.* **32**, 36–52.
- Whiten, A., Horner, V. & de Waal, F. 2005 Conformity to cultural norms of tool use in chimpanzees. *Nature* **437**, 737–740. (doi:10.1038/nature04047)
- Whiten, A., Spiteri, A., Horner, V., Bonnie, K. E., Lambeth, S. P., Schapiro, S. J. & de Waal, F. B. M. 2007 Transmission of multiple traditions within and between chimpanzee groups. *Curr. Biol.* **17**, 1038–1043. (doi:10.1016/j.cub.2007.05.031)
- Zajonc, R. B. 1965 Social facilitation. *Science* **149**, 269–274. (doi:10.1126/science.149.3681.269)