Evolution and Human Behavior 28 (2007) 11 – 17

Learning, productivity, and noise: an experimental study of cultural transmission on the Bolivian Altiplano☆

Charles Efferson[a,b,c,*], Peter J. Richerson[b,c,d], Richard McElreath[b,d,e,f], Mark Lubell[b,c], Ed Edsten[c,d], Timothy M. Waring[b,c], Brian Paciotti[b,c], William Baum[c]

[a]Santa Fe Institute, USA
[b]Graduate Group in Ecology, University of California, Davis, CA, USA
[c]Environmental Science and Policy, University of California, Davis, CA, USA
[d]Graduate Group in Animal Behavior, University of California, Davis, CA, USA
[e]Anthropology Department, University of California, Davis, CA, USA
[f]Graduate Group in Population Biology, University of California, Davis, CA, USA

Initial receipt 22 September 2005; final revision received 15 May 2006

Abstract

The theory of cultural transmission distinguishes between biased and unbiased social learning. Biases simply mean that social learning is not completely random. The distinction is critical because biases produce effects at the aggregate level that then feed back to influence individual behavior. This study presents an economic experiment designed specifically to see if players use social information in a biased way. The experiment was conducted among a group of subsistence pastoralists in southern Bolivia. Treatments were designed to test for two widely discussed forms of biased social learning: a tendency to imitate success and a tendency to follow the majority. The analysis, based primarily on fitting specific evolutionary models to the data using maximum likelihood, found neither a clear tendency to imitate success nor conformity. Players instead seemed to rely largely on private feedback about their own personal histories of choices and payoffs. Nonetheless, improved performance in one treatment provides evidence for some important but currently unspecified social effect. Given existing experimental work on cultural transmission from other societies, the current study suggests that social learning is potentially conditional and culturally specific.
© 2007 Elsevier Inc. All rights reserved.

Keywords: Cultural evolution; Nonlinear dynamics; Experimental games

1. Introduction

Imagine a population living in an environment with two productive technologies. Both technologies are generally available and involve equal explicit costs. As an apt example, consider a group of subsistence pastoralists facing the option of herding either sheep or llamas. Any individual can use either technology during any period. Livestock yields have a random component, however, and individuals do not know which technology brings the highest average yield. Thus, on average, choosing suboptimally involves an opportunity cost. At one extreme, our featured pastoralists could simply flip a coin periodically to determine which animals to raise. Although a straightforward approach, short of actually biasing choices toward the suboptimal technology, it would minimize the productivity of each individual and, by extension, the economy. At the other extreme, each pastoralist could have a genetic variant coding for a strong desire to raise the livestock variety that happens to be optimal in the native environment. In this case, everyone generally makes the best choice, with or without experience, so long as the environment and set of available technologies remain constant. If some of our pastoralists decide to pick up and move, however, or if weather patterns change over time, the resulting environmental changes might render the old optimum suboptimal. Similarly, if some process of innovation or technology transfer introduces a new

☆ This research was funded by a Jastro Shields Award through the Graduate Group in Ecology, University of California, Davis.
* Corresponding author. 1399 Hyde Park Road, Santa Fe, NM 87501.
E-mail address: cmefferson@santafe.edu (C. Efferson).

1090-5138/07/$ – see front matter © 2007 Elsevier Inc. All rights reserved.
doi:10.1016/j.evolhumbehav.2006.05.005

domesticate into the population, old technologies could become obsolete. In either case, genetic evolution may not be able to track the optimum over relevant time scales.

Between the extremes of randomization and genetic encoding lies a continuum of learning strategies. On the one hand, each individual could ignore her peers, experiment with both technologies individually, and rely exclusively on trial and error to estimate which option is best. In such a society, learned behavior depends only on private information, and each person effectively duplicates the efforts of every other. On the other hand, individuals could observe the behaviors and even livestock yields of their peers and could incorporate this information into their own decisions. The social information available and the way people use it will affect the actual distribution of behaviors in the population through time.

A key question concerns whether people discriminate in some way when they use social information. For example, if every individual indiscriminately picked another random person and blindly copied that person's behavior, the distribution of behaviors in a large population would not change through time. In contrast, if every individual copied the behavior of a single individual who recently had an especially large payoff, the population would rapidly become fixed on this behavior. Both of these social learning rules are overly simplistic, but they illustrate the key difference between social learning that discriminates in some way and social learning that does not. The distinction is equivalent to the difference between biased and unbiased social learning examined by Boyd and Richerson (1985) and Richerson and Boyd (2005).

As a hypothesis about behavioral dynamics, biased social learning rests on weak assumptions. It depends on only two propositions: humans learn from other humans in ways that affect behavior, and they do not do so in a completely random fashion. Both of these propositions seem like truisms. Nonetheless, the nature of social learning is properly an empirical question, as is the nature of any interaction between individual learning and social learning. More to the point, simply concluding that social learning is not random leaves many possibilities on the table, and the differences matter. To sort among the many conceivable forms biased social learning could take, empiricism is essential.

Here we present an experimental economic study designed to study how people learn individually and socially. Our study intersects with the subject matter of many disciplines, including the study of economic growth and its focus on innovation versus imitation (Aghion, Harris, Howitt, & Vickers, 2001; Barro & Sala-I-Martin, 2004), the sociological study of the diffusion of innovations (Henrich, 2001; Rogers, 1995), the study of learning in economics (Camerer & Ho, 1999; Camerer, 2003; Fudenberg & Levine, 1998; Merlo & Schotter, 2003; Schlag, 1998, 1999), evolutionary game theory and its pervasive assumption of payoff-biased imitation (Bowles, 2004; Gintis, 2000), the evolutionary theory of cultural transmission in humans (Boyd & Richerson, 1985, 2005; Cavalli-Sforza & Feldman, 1981; Henrich & McElreath, 2003; Richerson & Boyd, 2005), and the empirical study of social learning in animal behavior (Dall, Giraldeau, Olsson, McNamara, & Stephens, 2005; Fragaszy & Perry, 2003; van Schaik et al., 2003). The experimental setting is the scenario described above with two technologies and stochastic payoffs. The task is to test for two biases that have figured prominently in the literature on social learning: a tendency to imitate successful individuals (Henrich & Gil-White, 2001; Offerman & Sonnemans, 1998; Offerman & Schotter, 2005) and a conformist tendency to adopt the most common behavior in the population (Boyd & Richerson, 1982; Henrich, 2001; Henrich & Boyd, 1998).

2. Experimental methods

We conducted all experiments in September 2004 in seven communities in the high-altitude zone (ca. 3600–3800 m) of the Sama Biological Reserve in southern Bolivia. The estimated population of the entire reserve, in both the low-altitude and the high-altitude zones, is 5500 people (Bluske, 2004). In this study population, the basic economic activity is subsistence herding. Sheep are traditional, but llamas are rapidly growing in popularity. As a consequence, the basic choice of how to make a living in this area is analogous to our experiment. Two technologies are available, payoffs are stochastic, and people are not certain which technology is best on average. Although the two technologies in our experiment were simply "red" versus "green," the analogy with the subsistence economy in this study population gives our experiment a high degree of external validity. Moreover, the choice of a nonstandard subject pool stems from a belief that behavioral experiments need to move beyond the confines of the standard university subject pool, as strongly suggested by recent cross-cultural studies in experimental economics (Henrich et al., 2001, 2004, 2005). Recent cultural evolution studies provide results from experiments similar to this one but conducted among university undergraduates in the United States (Baum, Richerson, Efferson, & Paciotti, 2004; McElreath et al., 2005) and Switzerland (Efferson, C., Lalive, R., Richerson, P. J., McElreath, R., & Lubell, M., unpublished data).

2.1. Technologies and payoff information

The two technologies in this experiment were "red" and "green" in the form of red and green index cards. On each card was a payoff in Bolivian centavos drawn from one of two truncated normal distributions with means 30 and 39. The color with mean 30 was suboptimal, while the color with mean 39 was optimal. The untruncated normal distributions had an S.D. of 12. Payoffs were truncated 2.5 S.D. above and below the means to prevent negative numbers from the lower tail of the suboptimal distribution. Truncation reduced the standard deviation of payoff distributions to slightly less than 12. Payoffs were also rounded to take integer values. Thus, the set of possible payoffs for the suboptimal color included the integers from 0 to 60, while the corresponding set for the optimal color included the integers from 9 to 69.

Because many people who live in Sama cannot read, we had a rubber stamp of a sheep made as a means of providing additional payoff information. The total range of possible payoffs was from 0 to 69. If a particular card had a payoff from the lower half of this interval (i.e., the integers 0–34), it also had one sheep stamp regardless of color. If the card had a payoff from the upper half (i.e., the integers 35–69), it had two sheep stamps. Thus, participants who could not read had the option of reducing the problem to identifying the color that brought two sheep stamps at the highest rate.

2.2. General procedures

Upon arriving in a community, we randomly assigned participants to either an individual learning treatment or one of two social learning treatments described below. Two stacks of cards, one green and the other red, were placed face down in front of each participant. Players were told that each card had a payoff in centavos, and that either color could bring high or low payoffs because the payoffs were "of luck." They were told that one of the two colors was better in the sense that it would bring high payoffs more often than the other. Players were told that, for those who could not read, the cards also had sheep stamps, with one sheep stamp indicating a small quantity of money and with two stamps indicating a large quantity of money. We told them that they would make 50 choices and, at the end of the experiment, they would be paid in real money by summing individually over all choices.

For each choice, the card from the top of the relevant stack was turned over and placed upright on top of any chosen cards of the same color from previous periods. After the experiment, players responded to a brief questionnaire individually and were then paid. Because Bolivian currency is difficult to obtain in large quantities in denominations of less than 50 centavos, total payoffs were rounded up to the nearest 50-centavo increment. These various numbers were chosen, in part, to ensure that the average total payoff approximated 20 Bolivianos, which was the area's going rate at the time for a day's worth of unskilled labor. The following describes how procedures varied by treatment. Table 1 summarizes the information available to players in each treatment.

Table 1
A summary of the information available in the three experimental treatments

Treatment	Private feedback	Social feedback ($t \geq 2$)
Individual	Realized payoff	None
Best color	Realized payoff	Color with highest payoff
Total distribution	Realized payoff	Distribution of colors

Social information was only provided in the two relevant treatments in Periods 2–50.

2.3. Individual treatment

Participants were in a private room without other participants. They communicated their choices to the experimenter either verbally or via gestures. After each stated choice, the experimenter turned the appropriate card over for the participant to see. With the exception of one blind man who participated in this treatment, the experimenter did not announce the payoff. The optimal color and the color on the player's right side were randomized over individuals.

2.4. Best-color treatment

A given group entered a private room, usually in the community school building or church, where the experiment would take place. Each individual sat in front of a large cardboard box containing her own two stacks of cards. In addition to the individual payoff information discussed above, players were also told that, at the end of each period, they would learn the color chosen by the player who had received the highest payoff that period. To conduct the experiment, two experimenters encircled the room asking each individual which color she wanted to choose. The individual reached inside the box without speaking and pointed. One of the experimenters turned the appropriate card over inside the box for the participant to observe. The other experimenter recorded the color chosen and the payoff. After all players had made a choice in a given period, the data-recording experimenter announced the color that had produced the highest payoff for that period based on the centavo information contained on each card. The experimenter did not announce the associated payoff, and so players only observed their own individual payoffs. The concept of a period was easily conveyed with the Spanish word "vuelta," which implies one pass around the room. Although the possibility was not discussed up front, a handful of periods produced two-color ties for the highest payoff. When this outcome occurred, the data-recording experimenter simply explained that two players had received the same highest payoff with one choosing red and the other choosing green. Ties for the highest payoff involving the same color were never announced, although, theoretically, two greens tying for the highest payoff is not the same situation as one green producing the best payoff. Individuals were asked to remain silent during the experiment. The optimal color was randomized over experimental groups, and the color on the right side within each box was randomized over individuals. All individuals within a group, however, had the same optimal color, and this fact was explained at the beginning of each experimental session. This treatment included a group of 4, a group of 5, a group of 6, and four groups of 11.

2.5. Total distribution treatment

This treatment was identical to the best-color treatment with two exceptions. First, after each period, the experimenter announced the number in the group choosing green and the number choosing red. Second, this treatment involved five groups of 11.

3. Basic results, statistical methods, dynamics, and model selection

For each player, we calculated a mean payoff per period. These quantities differ by experimental treatment [analysis of variance: $F(2, 159)=5.491$, $p=.005$]. Fig. 1 shows that a higher proportion of players typically chose optimally under the total distribution treatment as compared to individual and best-color treatments.

Table 2 records the mean payoff per period with 95% confidence intervals (95% CIs).

As Table 3 shows, the mean payoff per period in the total distribution treatment is significantly greater than those in both the individual and best-color treatments, while the same quantities for the individual learning and best-color treatments are not significantly different.

To examine the structure of this difference in choices and payoffs, with a particular focus on how the difference relates to social information, we developed an a priori set of dynamic models to fit to each of the three data sets. Each model makes specific assumptions about how changing information, both private and social, affects the choices players make through time. Given a set of such models, the task is to identify the model that fits the data best. This approach is well established in the experimental economics literature on learning, and Camerer (2003) provides a broad overview. The question is fundamentally empirical in the sense that the researcher seeks a model that summarizes the observed data better than some set of alternate models (Camerer & Ho, 1999).

With a few exceptions, the models we used were theoretical models of individual learning from the economics literature (Camerer & Ho, 1999; Fudenberg & Levine, 1998) and social learning models from the cultural evolutionary literature (Boyd & Richerson, 1985). The basic approach is to take a model developed for theoretical purposes, introduce a noise structure, and fit the model to the data using maximum likelihood (Efferson & Richerson, 2006; McElreath et al., 2005). Appendix A presents a detailed description of how we accomplished this for the present study. Here we simply describe the models we fit to a particular data set and summarize the results qualitatively. First, however, we briefly discuss how we compared alternate models.

To compare alternate models, we fit all models using maximum likelihood and, as a model selection criterion, we used a derivative form (AIC_c) of Akaike (1973) criterion, discussed in Burnham and Anderson (2002, p. 66). The basic idea behind the Akaike criterion is the following. Some process generated the data at hand. We do not know this process, but based on past research, we have various candidate models that we think represent this process more or less well. Because models are models, however, and not reality, summarizing reality with a model always results in loss of information. The Akaike criterion selects the model from a specified set of models that is estimated to lose the least amount of information (Burnham & Anderson, 2002). In this regard, it is an "information-theoretic criterion." In practical terms, it selects the model that yields the best maximum likelihood value, but it includes a penalty for adding parameters that one has to estimate with the same amount of data. Thus, in loose terms, the question is not "Can we improve the fit with a more complex model?" but rather "Can we improve the fit per unit of added complexity?" Burnham and Anderson (2002), Forster and Sober (1994), and Hilborn and Mangel (1997) provide extensive overviews of information-theoretic criteria like the Akaike criterion and their many advantages over more conventional

Table 2
Mean individual payoff (with 95% CI), by period, for each of the three treatments

Treatment	Sample size	95% CI
Individual	48	34.952±0.503
Best color	59	35.386±0.740
Total distribution	55	36.498±0.698

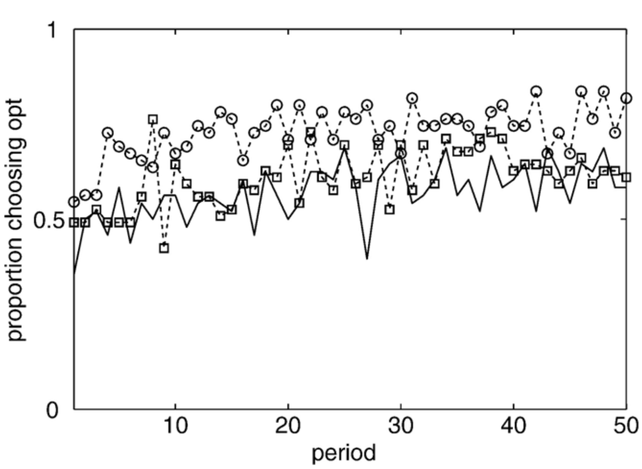

Fig. 1. The proportion of players choosing the optimal color, by period, for each of the three treatments. The solid line represents individual treatment, the dashed line with squares represents best-color treatment, and the dashed line with circles represents total distribution treatment.

Table 3
Multiple comparisons (Tukey–Kramer) of individual mean payoff, by period, in all combinations of two treatments

Treatment 1	Treatment 2	Lower bound	Estimate	Upper bound
Individual	Best color	−1.559	−0.434	0.690
Best color	Total distribution	−2.688	−1.545	−0.402
Total distribution	Total distribution	−2.196	−1.111	−0.026

The estimate indicated is the estimated difference in the individual mean payoff under Treatment 1 minus the same quantity under Treatment 2. The lower and upper bounds for this estimate at $\alpha=.05$ are also shown.

approaches to data analysis rooted in hypothesis testing. Efferson and Richerson (2006) discuss the use of the Akaike criterion specifically with respect to the experimental study of social learning.

4. Results

4.1. Results for individual learning treatment

We fit a total of six models to the individual learning data set. The simplest of these involved a constant probability of choosing the optimum in every period over all players. We also fit three models that do not explicitly incorporate individual histories of choices and payoffs but allow for a phenomenological description of trends and autocorrelated choices. Lastly, we fit two models that explicitly incorporate each individual's particular history of choices and payoffs. These models are called "attraction" models because they assume that the more money a player has made by choosing a particular option in the past, the stronger will be her attraction to it in the future. We describe the models and results in detail in Appendix B.

Based on this model-fitting exercise, players in this treatment did learn individually, although they tended to switch colors frequently from one period to the next. Learning can be seen in Fig. 1 because the proportion of individual learners choosing optimally increased steadily through time. The strong tendency to switch behaviors from one period to the next, however, prevented attraction models from fitting well. Attraction models incorporate individual choice and payoff histories and thus predict that a color generally producing higher payoffs (i.e., the optimal color in this experiment) will typically have a stronger and stronger attraction as time passes. In essence, players learned in this treatment, but the tendency to switch colors outweighed information processing as we model it with two standard attraction models from the economics literature (Camerer, 2003).

4.2. Results for best-color treatment

We fit the same individual learning models described above to this data set. We additionally used these individual learning models as a basis for models that combine individual learning and social learning. Social learning in this case takes the form of imitating the color that produced the highest payoff in the social group in the previous period. This model-fitting strategy is based on the following reasoning. The models combining individual and social learning often involve additional parameters beyond pure individual learning models. If the social learning component of the combined models does not sufficiently summarize systematic features of the data, the Akaike criterion will penalize these models for estimating relatively useless additional parameters associated with social learning. In such a case, pure individual learning models will fit better. If a combined model fits the best, it would mean that imitating success as we model it in Appendix A is a form of social learning that summarizes important features of the data.

As we explain in Appendix A, the two attraction models without social learning best summarize the data for the best-color treatment. This result has two implications. First, unless our models of imitating the best color from the previous period are completely inappropriate, players in this treatment apparently ignored the valuable social information provided to them. Social information was valuable in the sense that the announced best color in a given period was usually the color that was actually optimal for the session in question. Because payoffs were random, on occasion, the color that produced the highest payoff in a given period would be the color with the lower expectation (i.e., the suboptimal color), but this outcome was relatively uncommon. Appendix A includes a model that explains why this will generally be true. Second, although players in this treatment apparently relied on individual feedback, they did not learn individually in the same way as players in the individual learning treatment. In the best-color treatment, attraction models fit better than the other individual learning models. In the individual treatment, attraction models fit much worse relative to simple phenomenological models that do not account for individual choice and payoff histories. This result implies some kind of interaction between the social context and how players learned as individuals.

4.3. Results for total distribution treatment

The results for the total distribution treatment are analogous. In this case, to develop combined models of individual and social learning, we worked with three different models of frequency-dependent social learning: unbiased social learning and two different forms of conformity. Appendix A describes these models and Appendix B describes the results of the analysis in detail. As in the best-color treatment, the two attraction models without any social learning fit the best. Once again, this result suggests that players did not consistently use the social information provided in this treatment in any way captured by our models. Moreover, this information was valuable, as in the best-color case, because the color in the majority in any given period was almost always the color that was optimal for that session. This is because players in the total distribution treatment were learning individually. As a consequence, most of the players in a group in a particular period chose the color that was actually optimal, and so the group-level information we provided exaggerated the effects of individual learning into a valuable social signal. In short, conformity would have been an effective approach to making money.

4.4. Paradox and a posteriori analysis

The results above introduce a paradox. Relative to the individual treatment, the best-color treatment did not produce an effect in terms of optimality and average payoffs, but the total distribution treatment did. Learning in the best-color treatment as pure individual learners who ignored social information (an interpretation that the model-fitting exercise supports) could explain the lack of effect in the best-color treatment. The model-fitting exercise, however, also suggests that players in the total distribution treatment ignored social information and essentially learned as pure individual learners. What then is responsible for the observed increase in optimality and average payoffs in this latter treatment?

One possibility is a type of interaction psychologists refer to as social facilitation (Galef, 1988). Social facilitation is a form of learning in which learners do not acquire or use information from others. Rather, the presence of others encourages more efficient individual learning. For example, being in a group might cue individuals in a way that leads them to feel that they need to do well at the task. They might, as a consequence, try harder to estimate which color is optimal. Such an effect might explain why participants in the total distribution treatment performed better than those in the individual treatment, despite so little evidence that they used the information embedded in the social signal. In support of this interpretation, the model-fitting exercise indicates that players learned differently based on their private feedback in the individual treatment, where attraction models did not fit well, and based on their private feedback in the total distribution treatment, where attraction models did fit well. This difference seems especially salient in light of the fact that attraction models represent a more sophisticated use of private feedback than the other models under consideration. Nonetheless, if social facilitation is the only relevant mechanism behind the difference between the individual and the total distribution treatments, why then did placing players in groups in the best-color treatment not have the same effect? The absence of an effect here suggests that an additional mechanism is at play in the total distribution treatment.

In particular, Fig. 1 shows that the increase in optimal choices in the total distribution treatment seems to have come only in the first six or seven periods. In these periods, the proportion of players choosing the optimal technology rose quickly, relative to the other two treatments, but subsequently, there is no obvious difference in the trends found in the three treatments. To investigate this, we fit all of our models for the total distribution data set to the first seven periods only. Two results came out of the analysis. First, attraction models no longer fit the best, but rather the same phenomenological models that fit the individual treatment best also best fit the data from the first seven periods of the total distribution treatment. Second, the results suggested that players may have been using conformity in the first seven periods, but the limited amount of data produced confidence intervals too large to draw any firm conclusions for or against a social learning effect. On a related note, McElreath et al. (2005) conducted a similar experiment and found a strong interest in social information in early periods that waned rapidly as the experiment progressed.

5. Discussion and conclusion

Some simple facts best illustrate the value of the group-level information provided in the two social treatments. In the best-color treatment, across seven experimental sessions, the announced highest-paying color was the true optimum 274 of 350 times. Even more dramatically, in the total distribution treatment, a majority of players in the group chose the optimal color 237 of 250 times. Thus, in both experiments, social information aggregated a lot of individual noise into a valuable signal that reliably pointed toward the optimal color. Yet, on balance, model fitting provided little evidence that players used this signal. The one possible exception involved the initial periods of the total distribution treatment. Our analysis strongly suggests that information use was distinctive in these periods, and this difference accounts for the higher average payoffs for players in this treatment.

What we can say is that, in the first few periods of the total distribution treatment, some kind of interaction between individual learning and social situation produced an increase in optimal choices with a corresponding increase in productivity. The net effect over all 50 periods was modest because the increase in optimality was confined to the early stages of the experiment. The difference in absolute productivity, however, could be substantial in a real-world setting involving aggregation over hundreds or thousands of individuals and multiple time periods.

In the final analysis, we suspect that variation in social learning will have a kind of metastructure that will require a sustained experimental program to unpack. The current experiments found an interaction between social setting and behavior that affected payoffs, but the effect was not readily attributable to biased social learning. In contrast, Baum et al. (2004), Efferson et al. (unpublished data), and McElreath et al. (2005) found considerable evidence for cultural traditions in the laboratory and biased social learning. This kind of variation across studies could be part of a larger structure governing social learning. After many years of economic experiments on social preferences, for example, we now have some sense of how these preferences and their aggregate effects depend on and interact with the institutional setting, culture, and framing (Camerer, 2003; Camerer & Fehr, 2006; Fehr & Fischbacher, 2003, 2004; Henrich et al., 2004). We anticipate that the experimental study of social learning will prove to be similarly rich with subtleties.

Acknowledgments

C.E. thanks Gonzalo Reyes, Isabel López, the residents of Sama, the U.S. Environmental Protection Agency (fellowship no. U-91615601-3), and the Institute for Empirical Research in Economics, University of Zurich.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at doi:10.1016/j.evolhumbehav.2006.05.005.

References

Aghion, P., Harris, C., Howitt, P., & Vickers, J. (2001). Competition, imitation and growth with step-by-step innovation. *Review of Economic Studies, 68* (3), 467–92.

Akaike, H. (1973). Information theory as an extension of the maximum likelihood principle. In B. N. Petrov & F. Csaki (Eds.), *Second international symposium on information theory* (pp. 267–281). Budapest: Akademiai Kiado.

Barro, R. J., & Sala-I-Martin, X. (2004). *Economic growth*. (2nd ed). Cambridge: The MIT Press.

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* (5), 305–26.

Bluske, R. A. (2004). *Áreas Protegidas del Departamento de Tarija*. (2nd ed). Tarija, Bolivia: Protección del Medio Ambiente Tarija.

Bowles, S. (2004). *Microeconomics: Behavior, institutions, and evolution*. New York: Russell Sage.

Boyd, R., & Richerson, P. J. (1982). Cultural transmission and the evolution of cooperative behavior. *Human Ecology, 10* (3), 325–51.

Boyd, R., & Richerson, P. J. (1985). *Culture and the evolutionary process*. Chicago: University of Chicago Press.

Boyd, R., & Richerson, P.J. (2005). *The origin and evolution of cultures*. Oxford: Oxford University Press.

Burnham, K. P., & Anderson, D. R. (2002). *Model selection and multimodel inference: A practical information-theoretic approach*. (2nd ed). New York: Springer-Verlag.

Camerer, C., & Ho, T.-H. (1999). Experience-weighted attraction in normal form games. *Econometrica, 67* (4), 827–74.

Camerer, C. F. (2003). *Behavioral game theory: Experiments in strategic interaction*. Princeton: Princeton University Press.

Camerer, C. F., & Fehr, E. (2006). When does "economic man" dominate social behavior? *Science, 311* (6), 47–52.

Cavalli-Sforza, L. L., & Feldman, M. W. (1981). *Cultural transmission and evolution: A quantitative approach*. Princeton: Princeton University Press.

Dall, S. R. X., Giraldeau, L.-A., Olsson, O., McNamara, J. M., & Stephens, D. W. (2005). Information and its use by animals in evolutionary ecology. *Trends in Ecology and Evolution, 20* (4), 187–93.

Efferson, C., & Richerson, P. J. (2006). A prolegomenon to nonlinear empiricism in the human behavioral sciences. *Biology and Philosophy,* 2006 in press.

Fehr, E., & Fischbacher, U. (2003). The nature of human altruism. *Nature, 425* (23), 785–91.

Fehr, E., & Fischbacher, U. (2004). Social norms and human cooperation. *Trends in Cognitive Sciences, 8* (4), 185–90.

Forster, M. R., & Sober, E. (1994). How to tell when simpler, more unified, or less ad hoc theories provide more accurate predictions. *British Journal for the Philosophy of Science,* 45, 1–35.

Fragaszy, D. M., & Perry, S. (2003). *The biology of traditions: Models and evidence*. Cambridge: Cambridge University Press.

Fudenberg, D., & Levine, D. K. (1998). *The theory of learning in games*. Cambridge: The MIT Press.

Galef, B. G. J. (1988). Imitation in animals: History, definition, and interpretations of data from the psychological laboratory. In T. Zentall, J. Bennett, & G. Galef (Eds.), *Social learning: Psychological and biological perspectives* (pp. 3–28). Hillsdale, NJ: Lawrence Erlbaum.

Gintis, H. (2000). *Game theory evolving*. Princeton: Princeton University Press.

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* (4), 215–41.

Henrich, J., Boyd, R., Bowles, S., Camerer, C., Fehr, E., Gintis, H., et al. (2001). In search of homo economicus: Behavioral experiments in 15 small-scale societies. *American Economic Review, 91* (2), 73–8.

Henrich, J., Boyd, R., Bowles, S., Camerer, C., Fehr, E., Gintis, H., et al. (2005). "Economic man" in cross-cultural perspective: Behavioral experiments in 15 small-scale societies. *Behavioral and Brain Sciences, 28* (6), 795–815.

Henrich, J., & Gil-White, F. (2001). The evolution of prestige: Freely conferred deference as a mechanism for enhancing the benefits of cultural transmission. *Evolution and Human Behavior, 22* (3), 165–96.

Henrich, J., & McElreath, R. (2003). The evolution of cultural evolution. *Evolutionary Anthropology, 12* (3), 123–35.

Henrich, J., Robert, B., Bowles, S., Camerer, C., Fehr, E., & Gintis, H. (2004). *Foundations of human sociality: Economic experiments and ethnographic evidence from fifteen small-scale societies*. Oxford: Oxford University Press.

Hilborn, R., & Mangel, M. (1997). *The ecological detective: Confronting models with data*. Princeton: Princeton University Press.

McElreath, R., Lubell, M., Richerson, P.J., Waring, T.M., Baum, W., Edsten, E., et al. (2005). Applying evolutionary models to the laboratory study of social learning. *Evolution and Human Behavior, 26* (6), 483–508.

Merlo, A., & Schotter, A. (2003). Learning by not doing: An experimental investigation of observational learning. *Games and Economic Behavior, 42* (1), 116–36.

Offerman, T., & Schotter, J. (2005). *Sampling for information or sampling for imitation: An experiment on social learning with information on ranks*. Mimeo: New York University.

Offerman, T., & Sonnemans, J. (1998). Learning by experience and learning by imitating successful others. *Journal of Economic Behavior and Organization, 34* (4), 559–75.

Richerson, P. J., & Boyd, R. (2005). *Not by genes alone: How culture transformed the evolutionary process*. Chicago: University of Chicago Press.

Rogers, E. M. (1995). *Diffusion of innovations*. (4th ed). New York: The Free Press.

Schlag, K. H. (1998). Why imitate, and if so, how? A boundedly rational approach to multi-armed bandits. *Journal of Economic Theory, 78* (1), 130–56.

Schlag, K. H. (1999). Which one should I imitate? *Journal of Mathematical Economics, 31* (4), 493–522.

van Schaik, C. P., Ancrenaz, M., Borgen, G., Galdikas, B., Knott, C. D., Singleton, I., et al (2003). Orangutan cultures and the evolution of material culture. *Science, 299* (3), 102–5.

A Description of Models

The following details the models we fit to the three data sets. In all cases the approach involves taking a dynamical model (e.g. a model lifted from the cultural evolutionary literature), introducing a noise structure, and deriving a joint probability distribution over a data set. Table 1 summarizes the various models fit to each of the three data sets.

A.1 General Framework for all Dynamical Models

Define the choice of player b in period t as $c_{b,t} \in \{0, 1\}$, where $c_{b,t} = 0$ denotes that the player chooses the sub-optimal color, and $c_{b,t} = 1$ signifies the optimum. Choices for an entire data set are $c \in \Re^B \times \Re^T$, where B is the number of individuals in a given treatment, and T is the number of periods. Each model below includes a vector of parameters, Ω, and specifies a probability that b chooses optimally in t. In all cases the log likelihood function is

$$\ln\{L(\Omega \mid c)\} = \sum_{b=1}^{B} \sum_{t=1}^{T} (c_{b,t}) \ln \left\{ \frac{P(c_{b,t} = 1)}{P(c_{b,t} = 0)} \right\} + \ln\{P(c_{b,t} = 0)\},$$

which is simply derived from a joint probability distribution over a data set, c, given that observations are Bernoulli random variables.

As discussed in the main text, we use a derivative form of Akaike's criterion (Akaike, 1973) discussed in Burnham and Anderson (2002, p. 66) as a model selection criterion. In particular, when discussing the results of our model fitting, we use the term "Akaike weight." Akaike weights are essentially rescalings of Akaike values that allow us to summarize the proportional weight of evidence for each model in a set of models that have been fit to the same data set. As proportional weights, Akaike weights sum to 1 over a given set of models, and a large weight indicates substantial support for the associated model relative to other models under consideration (Burnham and Anderson, 2002). Unlike hypothesis testing, however, Akaike weights involve no arbitrary threshold (e.g. $\alpha = 0.05$) separating one categorical conclusion from another (e.g. significant versus non-significant).

A.2 Binomial

The simplest model assumes choices are binomially distributed with a constant probability of choosing the optimal color such that, $\forall t \in \{1, \ldots, T\}$ and $\forall b \in \{1, \ldots, B\}$,

$$P(c_{b,t} = 1) = p. \tag{1}$$

Model (1) neither incorporates individual feedback, nor does it allow for any interesting dynamical properties like trends or cyclical behavior.

A.3 Switching

The switching model, $\forall b$, is the following:

$$\begin{aligned} t = 1 \qquad & P(c_{b,t} = 1) = \beta \\ \forall t \in \{2, \ldots, T\} \qquad & P(c_{b,t} = 1 \mid c_{b,t-1} = 0) = \gamma \\ \forall t \in \{2, \ldots, T\} \qquad & P(c_{b,t} = 1 \mid c_{b,t-1} = 1) = 1 - \theta. \end{aligned} \tag{2}$$

Model (2) ignores individual payoff information and predicts choices based on a memory of only one period. If players tend to switch from the sub-optimal color to the optimal color more than vice versa, the parameter estimates will satisfy $\gamma > \theta$.

A.4 Signal Models

Boyd and Richerson (1995), Henrich and Boyd (1998), and Henrich (2001) have developed a straightforward theoretical framework for integrating individual and social learning. Here we use this framework as a basis for data analysis and refer to this class of models as "signal" models. We fit two types of signal models to the individual learning data set. These models assume that in each period each player receives a private signal about the relative merits of the two colors. The value of this signal in any given period is a realization of a random variable. With probability P_0 the player interprets the signal as conclusively indicating that color 0 (i.e. the sub-optimal technology) is optimal. With probability P_1 the individual interprets the signal as conclusively indicating that color 1 (i.e. the optimal technology) is optimal. With probability L the individual views the signal as inconclusive (e.g. because the realization of the random variable does not exceed some critical threshold) and relies on some other kind of information to make a decision. Notice that if individual learning is effective, it should be true that $P_1 > P_0 = 1 - P_1 - L$, which simply means that individual learning biases choices toward the optimum more than it does toward the sub-optimum. In the "signal maintain" model, the individual maintains her choice of color from the previous period when the signal is inconclusive. The full model is

$$\begin{aligned} t = 1 \qquad & P(c_{b,t} = 1) = P_1 + L(1/2) \\ \forall t \in \{2, \ldots, T\} \qquad & P(c_{b,t} = 1) = P_1 + Lc_{b,t-1}. \end{aligned} \tag{3}$$

Model (3) makes no explicit use of the actual individual feedback in an experimental session, and it assumes individual choices exhibit some positive autocorrelation at lag 1. The "signal switch" model is similar, except it assumes that individuals switch behavior under an inconclusive signal. Specifically,

$$\begin{aligned} t = 1 \qquad & P(c_{b,t} = 1) = P_1 + L(1/2) \\ \forall t \in \{2, \ldots, T\} \qquad & P(c_{b,t} = 1) = P_1 + L(1 - c_{b,t-1}). \end{aligned} \tag{4}$$

A.5 Attraction Models

Compared to the models above, a more complex approach to individual learning is to base choice probabilities on the particular history of choices and payoffs a given individual has experienced. We fit two such models to the individual learning data set. In all cases, each color has a current attraction that, *ceteris paribus*, is increasing in the payoffs one has received by choosing that color in the past. Loosely, the more money you've made with a particular color in the past, the stronger your attraction to it in the future. Formally, define the attraction to the sub-optimal color by player b at the end of period $t-1$ (i.e. after the choice in $t-1$ and before the choice in t) as $A^0_{b,t-1}$. The analogous attraction to the optimum is $A^1_{b,t-1}$. Lastly, the realized payoff for player b in t is $\pi_{b,t}$.

Given these definitions, a model of average reinforcement posits that, $\forall t \in \{1, \ldots, T\}$ and $j \in \{0, 1\}$,

$$A^j_{b,t} = \phi A^j_{b,t-1} + I(c_{b,t}, j)(1 - \phi)\pi_{b,t}, \tag{5}$$

where $c_{b,t} = j \Rightarrow I(c_{b,t}, j) = 1$, and $c_{b,t} \neq j \Rightarrow I(c_{b,t}, j) = 0$. The other attraction model is the model of experience-weighted attraction developed by Camerer and Ho (1999), which subsumes average reinforcement as a special case. Experience-weighted attraction specifies that

$$\begin{aligned} N_t &= \rho N_{t-1} + 1 \\ A^j_{b,t} &= \frac{\phi N_{t-1} A^j_{b,t-1} + I(c_{b,t}, j)\pi_{b,t}}{N_t}. \end{aligned} \tag{6}$$

For both attraction models, we convert attraction values to choice probabilities using the logit transformation:

$$\forall t \in \{1, \ldots, T\} \quad P(c_{b,t} = 1) = \frac{\exp\left(\lambda A^1_{b,t-1}\right)}{\exp\left(\lambda A^0_{b,t-1}\right) + \exp\left(\lambda A^1_{b,t-1}\right)}. \tag{7}$$

As with all dynamical systems, we need initial conditions for the state variables. For our purposes these initial values are parameters to estimate from the data: A^0_{init}, A^1_{init}, and N_{init}. Under the logit transformation, however, one cannot uniquely identify both initial attractions. So in all cases we assume $A^0_{init} = 0$ and estimate A^1_{init}. In addition, following Camerer and Ho (1999) we apply the following restrictions when applicable: $\phi, \lambda \geq 0$, $\rho \in [0, 1]$, $A^1_{init} \in [0, \overline{\pi}]$, where $\overline{\pi}$ is the maximum possible payoff in a period, and $N_{init} \in [0, 1/(1 - \rho)]$.

A.6 Payoff-Dependent Social Learning Models: Imitate the Best

For any draw from the the sub-optimal payoff distribution, let the payoff be a random variable, X, with p.d.f. $f(x)$ and c.d.f. $F(x)$. Similarly, the payoff for any draw from the optimal distribution is Y distributed according to $g(y)$ and $G(y)$. Let $I_t \in \{0, 1, \ldots, N\}$ be a variable specifying the number of players in the social group choosing the optimal color in period t. This information is important in deriving a model to fit to the data, but it was not available to players in the experiment. If $i_t = N$, someone who imitates the best in $t+1$ chooses the optimal color with certainty and never chooses the sub-optimum. Analogously, if $i_t = 0$, someone who imitates the best always chooses the sub-optimal color.

For the moment, assume $0 < i_t < N$, and thus a mix of choices exists. To simplify the notation below, temporarily drop the "t" subscript from i_t, but note that in what follows $i = i_t$. Designate the $N - i$ draws from the sub-optimal distribution as $X_1, \ldots, X_{N-i}$ and their corresponding order statistics as $X_{(1)} \leq X_{(2)} \leq \ldots \leq X_{(N-i)}$. Similarly the i

draws from the optimal distribution are $Y_1, \ldots, Y_i$, and the resulting order statistics are $Y_{(1)} \leq Y_{(2)} \leq \ldots \leq Y_{(i)}$. Under a mix of choices the density[1] over the highest payoff is

$$f_{X_{(N-i)}}(x) = (N-i)\{F(x)\}^{N-i-1} f(x).$$

Importantly, our experiments involved truncated normal distributions. Consequently, the p.d.f., $f(x)$, and c.d.f., $F(x)$, must be defined appropriately for the above expressions to be correct. Specifically, let $u(x)$ be the density function for the untruncated normal distribution, $N(30, 12)$. The density for $f(x)$ is thus defined as

$$\begin{aligned} x < 0 \quad &\Rightarrow \quad f(x) = 0 \\ 0 \leq x \leq 60 \quad &\Rightarrow \quad f(x) = \frac{u(x)}{\int_0^{60} u(x)dx} \\ 60 < x \quad &\Rightarrow \quad f(x) = 0. \end{aligned} \tag{8}$$

Given $f(x)$, the function $F(x)$ is then a typical c.d.f. Similar logic applies to the truncated payoff distribution for the optimal color.

The density over the highest payoff from the distribution for the optimal color is

$$g_{Y_{(i)}}(y) = i\{G(y)\}^{i-1} g(y).$$

Consequently, the *ex ante* conditional probability the highest payoff in the sample comes from the optimal distribution, denoted $k_1(i)$, is

$$k_1(i) = P\left(X_{(N-i)} < Y_{(i)} \mid 0 < i < N\right) = 1 - \int_{-\infty}^{\infty} \int_{-\infty}^{x} f_{X_{(N-i)}}(x) g_{Y_{(i)}}(y)\, dy dx.$$

[1]Ross (2002, p. 275) provides a general derivation of order-statistic densities.

After picking up the "t" subscript again, the complete model specifies the probability, $K(i_t)$, that the highest paying color in a period is the actual optimum:

$$\begin{aligned} i_t = 0 \quad &\Rightarrow \quad K(i_t) = 0 \\ i_t \in \{1, \ldots, N-1\} \quad &\Rightarrow \quad K(i_t) = k_1(i_t) \\ i_t = N \quad &\Rightarrow \quad K(i_t) = 1. \end{aligned} \tag{9}$$

Figure 1 shows the probabilities generated by (9) under the payoff distributions used in the experiment and groups of size 11. As the Figure shows, imitating the best color is a positively biased heuristic in the sense that it biases choices toward the optimal color when compared to unbiased transmission, which is also shown in Figure 1 as a reference.

A.7 Frequency-Dependent Social Learning Models: Linear Transmission and Conformity

To develop a set of models that incorporate frequency-dependent social learning, we now derive models specifying the probability a focal individual chooses the optimal color given that she makes her decision based on social information as opposed to individual feedback. Let $S_{b,t}$ be this conditional probability for individual b in period t. Further define q_t as the frequency of individuals in the group choosing the optimal color in t. Unbiased social learning (also known as "linear transmission") stipulates that

$$S_{b,t} = q_{t-1}. \tag{10}$$

A simple model of conformity (Boyd and Richerson, 1985) for groups of 11 players, as in our experiment, is

$$S_{b,t} = q_t(1-D) + D\sum_{i_t=6}^{11}\binom{11}{i_t}(q_{t-1})^{i_t}(1-q_{t-1})^{11-i_t}, \tag{11}$$

where D is a parameter that measures the nature of frequency-dependence. To reduce the number of parameters we have to estimate from the data, we assume $D = 1$. As described in Efferson *et al.* (2006), the assumption $D = 1$ is not entirely appropriate for the present experiments. The resulting model, however, has all the essential features of conformity and does not require us to estimate D in a maximum-likelihood setting that is already highly non-linear. Estimating D would add complexity to the estimation procedure and provide little insight with respect to the question of whether players were conformists. Efferson *et al.* (2006) present results from a related experiment in which estimating D is practical. In that study, estimating D provides a quantitative measure of frequency-dependent social learning, but it does not change the general conclusions in any way.

To allow for a stronger form of conformity, we also use another model that assumes individuals exhibit a highly pronounced tendency to adopt the most common color:

$$S_{b,t} = Z_{b,t-1}, \tag{12}$$

where $q_{t-1} \in [0, 1/2) \Rightarrow Z_{b,t-1} = 0$, $q_{t-1} = 1/2 \Rightarrow Z_{b,t-1} = 1/2$, and $q_{t-1} \in (1/2, 1] \Rightarrow Z_{b,t-1} = 1$.

A.8 Joint Models of Individual and Social Learning

For the best-color data set, imitating the best based on the signal models above is

$$\begin{aligned} t = 1 \qquad & P(c_{b,t} = 1) = P_1 + L(1/2) \\ \forall t \in \{2, \dots, T\} \qquad & P(c_{b,t} = 1) = P_1 + LK(i_{t-1}). \end{aligned} \tag{13}$$

Additionally define $I^1_{b,t}$ as the probability that player b chooses the optimal technology in t under individual learning as modeled by either the switching model (2) or model (7) and one of the two attraction models (i.e. (5) or (6)). This approach provides three additional models that combine individual and social learning. All have the same generic structure,

$$\begin{aligned} t = 1 \qquad & P(c_{b,t} = 1) = I^1_{b,t} \\ \forall t \in \{2, \dots, T\} \qquad & P(c_{b,t} = 1) = (1 - \alpha) I^1_{b,t} + \alpha K(i_{t-1}), \end{aligned} \tag{14}$$

where $\alpha \in [0, 1]$ is a parameter that weights the reliance on individual versus social learning.

To develop a stronger form of imitating the best color, define the variable $V_t \in \{0, 1\}$ such that $V_t = 0$ if the highest paid color in t was the sub-optimal color for the experimental session in question, and $V_t = 1$ if the highest paid color in t was the optimum. This yields additional models of imitating the best. The first is a type of signal model,

$$\begin{aligned} t = 1 \qquad & P(c_{b,t} = 1) = P_1 + L(1/2) \\ \forall t \in \{2, \dots, T\} \qquad & P(c_{b,t} = 1) = P_1 + LV_{t-1}. \end{aligned} \tag{15}$$

The remaining models are again based on the switching model (2) or model (7) using one of the two attraction models,

$$\begin{aligned} t = 1 \qquad & P(c_{b,t} = 1) = I^1_{b,t} \\ \forall t \in \{2, \ldots, T\} \qquad & P(c_{b,t} = 1) = (1 - \alpha) I^1_{b,t} + \alpha V_{t-1}. \end{aligned} \tag{16}$$

For the total-distribution data sets, we adopt an analogous approach to models that include both individual and social learning. Linear transmission models take the form,

$$\begin{aligned} t = 1 \qquad & P(c_{b,t} = 1) = P_1 + L(1/2) \\ \forall t \in \{2, \ldots, T\} \qquad & P(c_{b,t} = 1) = P_1 + L q_{t-1}, \end{aligned} \tag{17}$$

or the following form under either the switching model or model (7) using one of the two attraction models:

$$\begin{aligned} t = 1 \qquad & P(c_{b,t} = 1) = I^1_{b,t} \\ \forall t \in \{2, \ldots, T\} \qquad & P(c_{b,t} = 1) = (1 - \alpha) I^1_{b,t} + \alpha q_{t-1}. \end{aligned} \tag{18}$$

Given conformity as specified by (11), a signal model when $D = 1$ and $N = 11$ is

$$\begin{aligned} t = 1 \qquad & P(c_{b,t} = 1) = P_1 + L(1/2) \\ \forall t \in \{2, \ldots, T\} \qquad & P(c_{b,t} = 1) = P_1 + L \left\{ \sum_{i_t = 6}^{11} \binom{11}{i_t} (q_{t-1})^{i_t} (1 - q_{t-1})^{11 - i_t} \right\}. \end{aligned} \tag{19}$$

An analogous set of models using either the switching model or model (7) and one of the two attraction models is

$$t = 1 \qquad P(c_{b,t} = 1) = I^1_{b,t} \tag{20}$$

$$\forall t \in \{2, \ldots, T\} \qquad P(c_{b,t} = 1) = (1 - \alpha) I^1_{b,t} + \alpha \left\{ \sum_{i_t=6}^{11} \binom{11}{i_t} (q_{t-1})^{i_t} (1 - q_{t-1})^{11-i_t} \right\}.$$

Lastly, the stronger form of conformity specified in (12) yields

$$t = 1 \qquad P(c_{b,t} = 1) = P_1 + L(1/2) \tag{21}$$

$$\forall t \in \{2, \ldots, T\} \qquad P(c_{b,t} = 1) = P_1 + L Z_{t-1}$$

and a set of models based on

$$t = 1 \qquad P(c_{b,t} = 1) = I^1_{b,t} \tag{22}$$

$$\forall t \in \{2, \ldots, T\} \qquad P(c_{b,t} = 1) = (1 - \alpha) I^1_{b,t} + \alpha Z_{t-1}.$$

B Detailed Results

Below we include detailed numerical results for those models that fit a particular data set well.

B.1 Results for Individual Learning Treatment

For the individual treatment, the switching model (2) fits the data best to an overwhelming degree. The only other model with a notable Akaike weight is the signal switch model (4). Table 2 shows the results for these two models. Because both of these models postulate a tendency to switch behavior from one period to the next, the good relative fit indicates that a dominant feature of play in the individual learning treatment was a tendency to switch technologies from one period to the next. As Figure 1 from the main text shows, however, individuals did learn in the sense that the trend through time is clearly toward optimality. This trend explains why the switching model fits the data better than the signal switch model. The switching model allows a distinction between switching to the optimum and switching away from the optimum, while the signal switch model does not. Indeed, the point estimate for the probability of switching to the optimum (0.7108) is greater than the estimate for the probability of switching away from the optimum (0.5295), and the approximate confidence intervals do not overlap (see Table 2). These results for the individual learning treatment indicate that players did learn, but their tendency to switch from one period to the next was sufficiently pronounced that individual learning models using particular payoff histories (i.e. the two attraction models) did not fit well in relative terms. In essence, the tendency to switch technology outweighed information processing as we model it with average reinforcement (5) and experience-weighted attraction (6).

B.2 Results for the Best-Color Treatment

Average reinforcement (5) and experience-weighted attraction (6), both without a social learning component, summarize the best-color data set best. Experience-weighted attraction (Akaike weight: 0.4433) fits better than average reinforcement (Akaike weight: 0.1196). Several models that combine average reinforcement or experience-weighted attraction with some form of imitating the best also generate substantial Akaike weights. Experience-weighted attraction mixed with the strong version of imitating the best, for example, has an Akaike weight of 0.1942. Tellingly, however, the estimated probability of using social information (α) is only 0.0122. The 95% confidence interval for this estimate overlaps 0, and thus we cannot unambiguously conclude that players in this population exhibit any tendency to imitate the best. Other versions of average reinforcement or experience-weighted attraction that include some version of imitating the best produce even smaller α estimates, suggesting that the individual learning components are essentially responsible for all the fit. As a consequence, on balance the analysis for the best-color data reveals that players relied almost exclusively on individual feedback to make decisions. Table 3 shows the results from the average reinforcement model and the experience-weighted attraction model, both without some form of imitating the best. Adding a tendency to imitate the best based on (7) produces similar results, but with the addition of extremely small estimated α values.

B.3 Results for the Total-Distribution Treatment

The results from the total-distribution treatment are analogous. Average reinforcement (5) and experience-weighted attraction (6), both without a social learning component, fit the best. These models, both with and without some form of social learning (i.e. linear, weak conformity, or strong conformity), account for virtually all of the total Akaike weight. All other models yield trivially small Akaike weights. As in the best-color case, any model that couples average reinforcement or experience-weighted attraction with some form of social learning produces extremely small estimated α values that cannot be distinguished from 0 based on the 95% confidence intervals. Once again, this result suggests that players did not consistently use the social information provided in this experiment in any way captured by our models. Unlike the best-color case, average reinforcement fits the total-distribution data better than experience-weighted attraction. All of the experience-weighted attraction models together yield a summed Akaike weight of approximately 0.118. Most of the remaining weight (i.e. $1 - 0.118 = 0.882$) goes to the various average reinforcement models, with those including some form of social learning all producing trivially small α estimates. Table 4 shows the results for the average reinforcement model without social learning.

B.4 Paradox and *A Posteriori* Analysis

As explained in the main text, these results are paradoxical. In particular, what explains the increase in average per-period payoffs in the total-distribution treatment? We investigated this by fitting all of our models for the total-distribution data set to the first 7 periods only. The models that fit the best were the signal switch model (4) and the switching models, including both the switching model by itself (2) and in combination with linear transmission (18), weak conformity (20), or strong conformity (22). Our model selection criterion did not distinguish clearly between the relative fits of these five models, suggesting that the data for the first 7 periods lend similar support to all of them. For the three social learning models, the estimated reliance on social learning (α) ranged from 0.0814 for the strong conformity model to 0.1930 for linear transmission. These point estimates imply that, unlike other cases described above, the constraint $\alpha \geq 0$ did not bind. Nonetheless, given the small amount of data when using only the first 7 periods, the confidence intervals were too large to draw any firm conclusions for or against a social learning effect.

References

Akaike, H. (1973). Information Theory as an Extension of the Maximum Likelihood Principle. In B. N. Petrov and F. Csaki, editors, *Second International Symposium on Information Theory*, pages 267–281. Budapest: Akademiai Kiado.

Boyd, R. and Richerson, P. J. (1985). *Culture and the Evolutionary Process*. Chicago: University of Chicago Press.

Boyd, R. and Richerson, P. J. (1995). Why Does Culture Increase Human Adaptability? *Ethology and Sociobiology*, **16**(2), 125–143.

Burnham, K. P. and Anderson, D. R. (2002). *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach*. New York: Springer-Verlag, 2nd edition.

Camerer, C. and Ho, T.-H. (1999). Experience-Weighted Attraction in Normal Form Games. *Econometrica*, **67**(4), 827–874.

Efferson, C., Lalive, R., Richerson, P. J., McElreath, R., and Lubell, M. (2006). Conformity and Indifference: The Structure of Frequency-Depedent Social Learning. Mimeo, University of Zuerich.

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. and Boyd, R. (1998). The Evolution of Conformist Transmission and the Emergence of Between-Group Differences. *Evolution and Human Behavior*, **19**(4), 215–241.

Ross, S. (2002). *A First Course in Probability*. Upper Saddle River, NJ: Prentice Hall, 6th edition.

Table 1: Models fit to each of the three data sets. Average reinforcement is designated "A.R." and experience-weighted attraction "E.W.A." In the cases involving an individual model with a social component (i.e. "w/ social"), all such individual models were coupled with all appropriate social models in all combinations.

Model	Individual	Best-Color	Total-Distribution
Binomial	X	X	X
Switching	X	X	X
Signal Maintain	X	X	X
Signal Switch	X	X	X
A.R.	X	X	X
E.W.A.	X	X	X
Switching w/ social		X	X
Signal w/ social		X	X
A.R. w/ social		X	X
E.W.A. w/ social		X	X
Weak imitate best		X	
Strong imitate best		X	
Linear			X
Weak conformity			X
Strong conformity			X

Table 2: Parameter estimates, maximized log likelihood, AIC_c, and Akaike weights (w_i) for the two models that best fit the 2400 choices from the individual treatment. Standard errors (S.E.) and confidence intervals (Lower CL and Upper CL) are derived from the Hessian of the negative log likelihood and are approximate.

Model	Parameter	Estimate	S.E.	Lower CL	Upper CL
	β	0.3541	0.0690	0.2188	0.4894
Switching	γ	0.7108	0.0142	0.6828	0.7384
	θ	0.5295	0.0136	0.5028	0.5562
$\ln L^*$:	-1566.2805	AIC_c:	3138.5709	w_i:	0.9870
Signal	P_1	0.4662	0.0136	0.4397	0.4928
Switch	L	0.2398	0.0198	0.2010	0.2787
$\ln L^*$:	-1571.6148	AIC_c:	3147.2346	w_i:	0.0130

Table 3: Parameter estimates, maximized log likelihood, AIC_c, and Akaike weights (w_i) for the average reinforcement model and the experience-weighted attraction model when fit to the 2950 choices from the best-color treatment. Standard errors (S.E.) and confidence intervals (Lower CL and Upper CL) are derived from the Hessian of the negative log likelihood and are approximate.

Model	Parameter	Estimate	S.E.	Lower CL	Upper CL
Average Reinforcement	ϕ	0.9358	0.0108	0.9146	0.9571
	λ	7.4587	0.6367	6.2109	8.7066
	A^0_{init}	0	N/A		
	A^1_{init}	0.0030	0.0131	-0.0226	0.0286
$\ln L^*$:	-1722.4398	AIC_c:	3450.8877	w_i:	0.1196
EWA	ϕ	0.9567	0.0127	0.9317	0.9816
	ρ	0.9675	0.0126	0.9429	0.9921
	λ	9.0643	1.1334	6.8429	11.2857
	A^0_{init}	0	N/A		
	A^1_{init}	3.8000×10^{-13}	0.0119	-0.0233	0.0233
	N_{init}	11.6008	1.1050	9.4350	13.7665
$\ln L^*$:	-1719.1237	AIC_c:	3448.2677	w_i:	0.4433

Table 4: Parameter estimates, maximized log likelihood, AIC$_c$, and Akaike weight (w_i) for the average reinforcement model when fit to the 2750 choices from the total-distribution treatment. Standard errors (S.E.) and confidence intervals (Lower CL and Upper CL) are derived from the Hessian of the negative log likelihood and are approximate.

Model	Parameter	Estimate	S.E.	Lower CL	Upper CL
Average Reinforcement	ϕ	0.9070	0.0112	0.8851	0.9289
	λ	7.2936	0.2768	6.7510	7.8362
	A^0_{init}	0	N/A		
	A^1_{init}	0.0643	0.0150	0.0348	0.0937
$\ln L^*$:	-1366.7368	AIC$_c$:	2739.4824	w_i:	0.4202

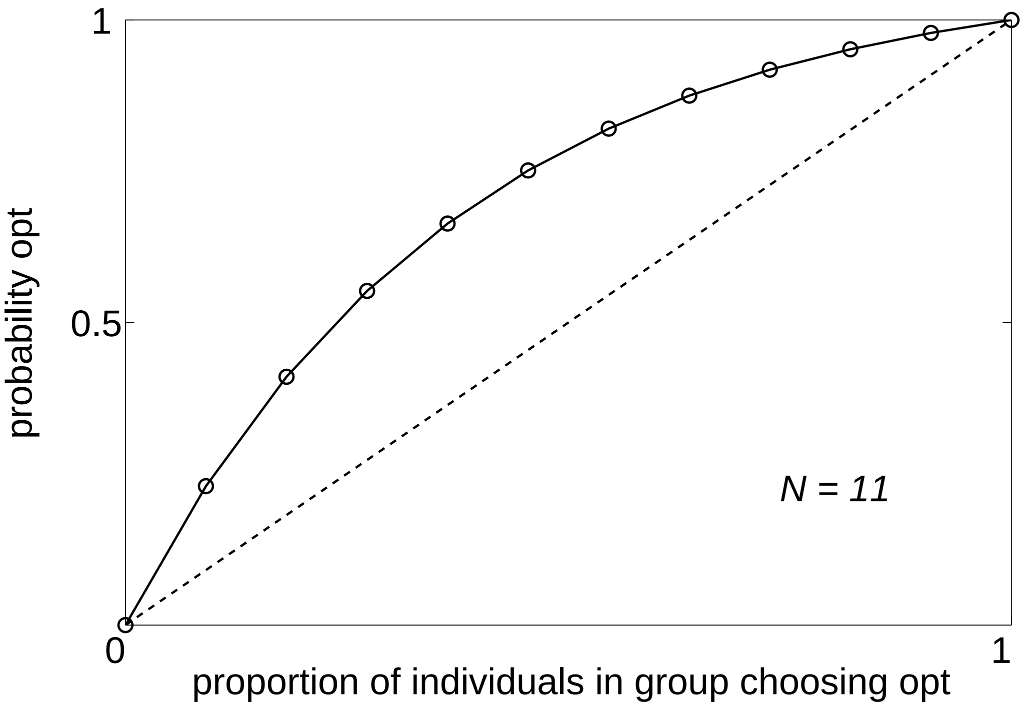

Figure 1: For groups of size 11 and payoff distributions used in the experiments, the probability that imitating the color producing the highest payoff will identify the optimal color. This function is shown as a line with circles. As a reference, unbiased transmission is the dashed line.