Conformists and Mavericks in the Lab: The Structure of Frequency-Dependent Social Learning*

Charles Efferson[1,2,4], Rafael Lalive[2,3], Peter J. Richerson[4,5], Richard McElreath[4,6], and Mark Lubell[4,5]

[1]Santa Fe Institute

[2]Institute for Empirical Research in Economics, University of Zürich

[3]Department of Econometrics and Political Economy, University of Lausanne

[4]Graduate Group in Ecology, University of California, Davis

[5]Environmental Science and Policy, University of California, Davis

[6]Anthropology, University of California, Davis

November 14, 2006

Running title: Frequency-dependent social learning

Word count: ≈ 1900

*__Corresponding author:__ Charles Efferson, Santa Fe Institute, 1399 Hyde Park Rd., Santa Fe, NM 87501, *cmefferson@santafe.edu.*

Abstract: Conformity is a type of social learning that has received considerable attention among social psychologists and human evolutionary ecologists, but existing empirical research does not identify conformity cleanly. Conformity is more than just a tendency to follow the majority; it involves an exaggerated tendency to follow the majority. We conducted an experiment to see if players were conformists by separating individual and social learners. Players chose between two technologies repeatedly. Payoffs were random, but one technology had a higher expected payoff. Individual learners knew their realized payoffs after each choice, while social learners only knew the distribution of choices among individual learners. A subset of social learners behaved according to a classic model of conformity. The remaining social learners did not respond to frequency information. Given this heterogeneity in learning strategies, a tendency to conform increased earnings dramatically.

Key words: cultural evolution, frequency-dependent social learning, experimental games

1 Introduction

Social scientists have long recognized the importance of frequency-dependent social learning (Asch, 1956; Boyd and Richerson, 1985; Henrich and McElreath, 2003; Bowles, 2004; Boyd and Richerson, 2005; Lumsden and Wilson, 1980; Richerson and Boyd, 2005; Sherif and Murphy, 1936). Frequency-dependent social learning postulates that individuals adopt a given behavior with a probability that varies as some function of how common the behavior is in a relevant social group. Conformity is a type of frequency dependence that has received considerable attention. In a simple case with two behaviors, R and B, where r_t is the frequency of R in the population, conformity means that in the near future an individual exhibits behavior R with a probability less than r_t if $r_t < 1/2$ but greater than r_t if $r_t > 1/2$. In other words, individuals have a *disproportionate* tendency to follow the majority. Theoretically, conformity is a valuable way to make good decisions in temporally and spatially variable environments (Henrich and Boyd, 1998).

Conformity also has important dynamical properties. It produces multiple steady states at the aggregate level and can lead otherwise similar societies to evolve in completely different ways in the wake of small random effects (Bowles, 2004; Efferson and Richerson, 2006; Salganik *et al.*, 2006). Like conformity, information cascades also lead populations to converge on different equilibria, but the mechanisms are entirely different from those we discuss. In information cascades, multiple equilibria are supported by a very particular interaction between sequential decisions and the limited accumulation of information at the group level (Anderson and Holt, 1997; Bikhchandani *et al.*, 1992). We focus instead on a pure conformist effect, by which we mean an exaggerated tendency to follow the majority.

Because conformity produces multiple stable steady states, it can reduce behavioral variation within social groups while increasing variation between social groups. Thus, in conjunction with the punishment of norm violations and the imitation of success, conformity plays a critical role in the study of how prosocial tendencies evolved in humans via cultural group selection (Boyd and Richerson, 1982; Boyd *et al.*, 2003; Fehr and Fischbacher, 2003, 2004; Fehr and Gaechter, 2002; Guererk *et al.*, 2006; Guzmán *et al.*, 2006; Henrich, 2004; Henrich and Boyd, 2001). Conformity also appears to be critical in explaining aggregate patterns that characterize the diffusion of technological innovations (Rogers, 1995; Henrich, 2001).

In spite of conformity's acknowledged importance, empirical research to date cannot identify conformity as a disproportionate tendency to follow the majority. Classic research in social psychology (Sherif and Murphy, 1936; Asch, 1956; Aronson *et al.*, 2002), for example, and recent experiments with chimpanzees (Whiten *et al.*, 2005) show that a focal individual is more likely to adopt a behavior as that behavior becomes more common. A simple model of non-conformity makes exactly the same prediction, as do other hypotheses about how individuals respond to frequency-dependent information (electronic supplement). Such distinctions are fundamental and not simply matters of definition. Conformity and non-conformity, for instance, predict radically different aggregate dynamics, with non-conformity increasing behavioral variation within groups and decreasing variation between groups (Efferson and Richerson, 2006). To demonstrate conformity, in addition to showing that individuals adopt common behaviors, researchers must show that this inclination is exaggerated in the way described above. Here we present a jointly theoretical and experimental approach to this problem.

The electronic supplement presents theoretical results formalizing how conformity works. As in the experiment described shortly, the theory focuses on situations in which the payoffs associated with different behaviors are stochastic. The best behavior is not obvious because feedback is noisy. In this case, conformity can filter noisy individual feedback into a powerful social signal that points clearly toward the best behavior. The effectiveness of conformity, however, depends on how good individual learning is. If individual learning is bad, conformity is worse. If individual learning is good, conformity is even better.

2 Experimental methods

With 70 students at the University of Zürich and the Swiss Federal Institute of Technology, we conducted the following experiment. In each period each player faced a choice between one of two technologies ("red" versus "blue"). Payoffs followed normal distributions, but one color was optimal in that its payoff distribution had a higher expectation. Players did not know which color was better. They made choices for six blocks of 25 periods each. Each block of 25 periods had a randomly selected optimal color, but all players who played together always had the same optimal color. All of this was explained in the instructions before beginning an experimental session. In addition, participants viewed an extensive demonstration before the beginning of the experiment. The demonstration produced various animated histograms that gave subjects an intuition for how the random payoffs were generated, even if they did not have formal training in probability theory. The electronic supplement provides more details.

Players within a session were divided into two groups that played simultaneously. In one group of 5 players, each player individually chose one of the two colors in each period and immediately received private information about her realized payoff. These players did not have any information about other players, and so we refer to them as *individual learners*. In the other group, composed of 6-7 players, each player had social information about the distribution of choices (e.g. 3 red, 2 blue) among the individual learners. These players did not have any information about their own payoffs, and thus we refer to them as *social learners*. Social information was available after all individual learners had made their choices in a period but before a given social learner had made her choice. After communicating the social information, each social learner made a choice between the two colors privately and received a payoff. Realized payoffs, however, were never communicated to players in this group, and individual learning was consequently not possible. The entire experiment was conducted on a local computer network using z-Tree (Fischbacher, 1999).

3 Results and Discussion

The value of conformity depends on the effectiveness of individual learning. As Figure 1 shows, individual learning was effective in that the proportion of individual learners choosing the optimal color increased substantially through time, and thus theory suggests that for social learners conformity should have been an effective approach to making money.

[Figure 1 about here]

To test for conformity, we take the formal definition of frequency-dependent social learning as initially formulated (Boyd and Richerson, 1982, 1985) and use it as our theoretical framework. If r_t is the frequency of individual learners choosing red in period t, the model stipulates for each social learner that $r_t < 1/2 \Rightarrow P(\text{social learner chooses red}) = r_t(1-D)$, $r_t = 1/2 \Rightarrow P(\text{social learner chooses red}) = 1/2$, and $r_t > 1/2 \Rightarrow P(\text{social learner chooses red}) = r_t + D(1 - r_t)$. The parameter $D \in [-1, 1]$ controls the nature of frequency dependence. When $D \in [-1, 0)$, social learning is non-conformist, when $D = 0$ social learning is linear (Boyd and Richerson, 1985), and when $D \in (0, 1]$ it is conformist.

To evaluate the theory, we estimated D using maximum likelihood under three different levels of assumed heterogeneity among social learners (electronic supplement). The simplest model posits a single value of D over all observations and all social learners. The second model divides social learners into two groups based on their response to a single questionnaire item. This questionnaire item allowed us, in an *a priori* fashion, to divide social learners into those who claimed they tended to follow the majority during the experiment and those who did not (see electronic supplement). We call these players respectively "stated conformists" and "not stated conformists," and our second model estimates a separate D value for each of these two groups. The final model estimates an individual value of D for each social learner (i.e. individual fixed effects). We use AIC_c, an improved form of Akaike's original criterion (Akaike, 1973; Burnham and Anderson, 2002), as a model-selection criterion.

Table 1 summarizes the results.

[Table 1 about here]

The model of individual fixed effects fits the best to an overwhelming degree, but the AIC_c values also indicate that the two-parameter model is a vast improvement over the simple model that estimates a single D value. This finding means that individual variation in frequency-dependent social learning is extremely important, but nonetheless the distinction between stated conformists and those who were not stated conformists also captures important systematic variation.

Figure 2 compares the data and the two-parameter version of the model.

[Figure 2 about here]

The model fits poorly for social learners who were not stated conformists but quite well for those who were stated conformists. The 12 social learners who were not stated conformists, in effect, did not respond on average to information about the frequencies of alternative behaviors in any notable way. Thus the model, though it can be fit using maximum likelihood, is not based on assumptions that were generally appropriate for social learners in this group. For the 28 stated conformists, however, data and model are nearly indistinguishable for much of the function's domain. The exceptions at the boundaries show that stated conformists had a small tendency to play the absent color when all five individual learners were choosing the same color.

Figure 3 additionally shows considerable individual variation in play among both stated conformists and social learners who were not stated conformists.

[Figure 3 about here]

Moreover, as predicted by theory (electronic supplement), a strong positive relationship between conformity and earnings exists. In short, because the individual learners were actually learning, social learners who showed a strong inclination to follow the majority among individual learners were the social learners who made the most money. Social learners who did not respond to the available frequency-dependent information left money on the table.

These results show that individual heterogeneity is critical to understanding frequency-dependent social learning. Specifically, our results suggest a meaningful distinction between those who conform and those who largely ignore information about behavioral frequencies. Nonetheless substantial individual variation also exists within each of these two generic groups of players. The implications of this kind of heterogeneity with respect to aggregate behavioral dynamics is a problem that has received no attention, but one obviously important consideration would involve how conformists and other types of learners assort into groups both within and between societies. For instance, if all else is equal and conformists do form groups assortatively, conformist groups should be more productive than their less conformist counterparts as long as some basis for effective individual learning is present. The current study documents considerable heterogeneity within a subject pool that should have a relatively high degree of cultural homogeneity. Given existing cross-cultural work in experimental economics (Henrich *et al.*, 2006), heterogeneity in social learning across societies could also be crucially important.

Acknowledgements

The authors would like to thank all the members in the chair of Ernst Fehr at the Institute for Empirical Research in Economics for assistance and comments. A particular thanks goes to Urs Fischbacher for his z-Tree program to produce animated histograms. CE thanks the U.S. Environmental Protection Agency (STAR Fellowship # U-91615601-3) for financial support. This research was funded by a U.S. National Science Foundation Grant (Award # 0340148 to PJR, RM, and ML).

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.

Anderson, L. R. and Holt, C. A. (1997). Information cascades in the laboratory. *The American Economic Review*, **87**(5), 847–862.

Aronson, E., Wilson, T. D., and Akert, R. M. (2002). *Social Psychology*. Upper Saddle River, NJ: Prentice-Hall, 4th edition.

Asch, S. E. (1956). Studies of independence and conformity: I. a minority of one against a unanimous majority. *Psychological Monographs*, **70**, 1–70.

Bikhchandani, S., Hirshleifer, D., and Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as information cascades. *The Journal of Political Economy*, **100**(5), 992–1026.

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

Boyd, R. and Richerson, P. J. (1982). Cultural transmission and the evolution of cooperative behavior. *Human Ecology*, **10**(3), 325–351.

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

Boyd, R. and Richerson, P. J. (2005). *The Origin and Evolution of Cultures*. Oxford: Oxford University Press.

Boyd, R., Gintis, H., Bowles, S., and Richerson, P. J. (2003). The evolution of altruistic punishment. *Proceedings of the National Academy of Sciences*, **100**(6), 3531–3535.

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

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

Fehr, E. and Fischbacher, U. (2003). The nature of human altruism. *Nature*, **425**(23 October), 785–791.

Fehr, E. and Fischbacher, U. (2004). Social norms and human cooperation. *Trends in Cognitive Sciences*, **8**(4), 185–190.

Fehr, E. and Gaechter, S. (2002). Altruistic punishment in humans. *Nature*, **415**(10 January), 137–140.

Fischbacher, U. (1999). z-Tree: toolbox for readymade economic experiments. IEW Working paper 21, University of Zuerich.

Guererk, O., Irlenbusch, B., and Rockenback, B. (2006). The competitive advantage of sanctioning institutions. *Science*, **312**(7 April), 108–111.

Guzmán, R. A., Rodríguez-Sickert, C., and Rowthorn, R. (2006). When in Rome, do as the Romans do: the coevolution of altruistic punishment, conformist learning, and cooperation. *Evolution and Human Behavior*, **in 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. (2004). Cultural group selection, coevolutionary processes and large-scale cooperation. *Journal of Economic Behavior and Organization*, **53**(1), 3–35.

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.

Henrich, J. and Boyd, R. (2001). Why punish defectors: weak conformist transmission can stabilize costly enforcement of norms in cooperative dilemmas. *Journal of Theoretical Biology*, **208**(1), 79–89.

Henrich, J. and McElreath, R. (2003). The evolution of cultural evolution. *Evolutionary Anthropology*, **12**(3), 123–135.

Henrich, J., McElreath, R., Barr, A., Ensminger, J., Barrett, C., Bolyanatz, A., Cardenas, J. C., Gurven, M., Gwako, E., Henrich, N., Lesorogol, C., Marlowe, F., Tracer, D., and Ziker, J. (2006). Costly punishment across human societies. *Science*, **312**(23 June), 1767–1770.

Lumsden, C. J. and Wilson, E. O. (1980). Translation of epigenetic rules of individual behavior into ethnographic patterns. *Proceedings of the National Academy of Sciences*, **77**(7), 4382–4386.

Newey, W. K. and West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, **55**(3), 703–708.

Richerson, P. J. and 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*. New York: The Free Press, 4th edition.

Salganik, M. J., Dodds, P. S., and Watts, D. J. (2006). Experimental study of inequality and unpredictability in an artificial cultural market. *Science*, **311**(10 February), 854–856.

Sherif, M. and Murphy, G. (1936). *The Psychology of Social Norms*. New York: Harper and Brothers.

Whiten, A., Horner, V., and de Waal, F. B. M. (2005). Conformity to cultural norms of tool use in chimpanzees. *Nature*, **437**(29 September), 737–740.

Table 1: Model description, the number of estimated parameters, maximized log likelihood ($\ln L^*$), Akaike value (AIC_c), and Akaike weight (w_i) for each of the three models fit to the social learners' data. Altogether the experiment produced 5000 observations for social learners, and 3749 of these could be used to estimate D and calculate AIC_c values (see electronic supplement). Akaike weights sum to 1 and summarize the proportional weight of evidence in support of each model, where larger weights indicate more support. The absolute difference between AIC_c values also has meaning (Burnham and Anderson, 2002) and is the basis for our claim that the two-parameter model is a vast improvement over the one-parameter model.

Model	Parameters	$\ln L^*$	AIC_c	w_i
Single D	1	−2159.28	4320.56	4.11×10^{-112}
Conformist, Y or N	2	−2014.39	4032.78	1.27×10^{-49}
Fixed Effects	40	−1863.36	3807.60	> 0.99

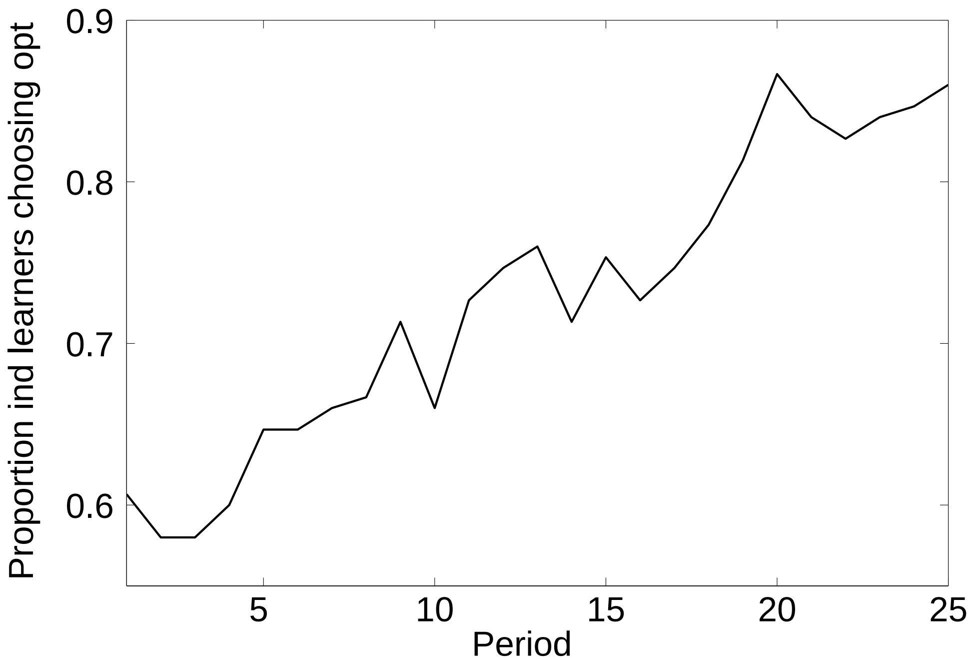

Figure 1: The proportion of individual learners choosing the optimal technology through time. The upward trend is highly significant ($p < 0.01$) when we regress the proportion choosing optimally on period using the method of Newey and West (1987) to correct for heteroskedasticity and autocorrelation up to lag 3. In this case, the estimated coefficient for period is 0.012 and the R^2 value is 0.930.

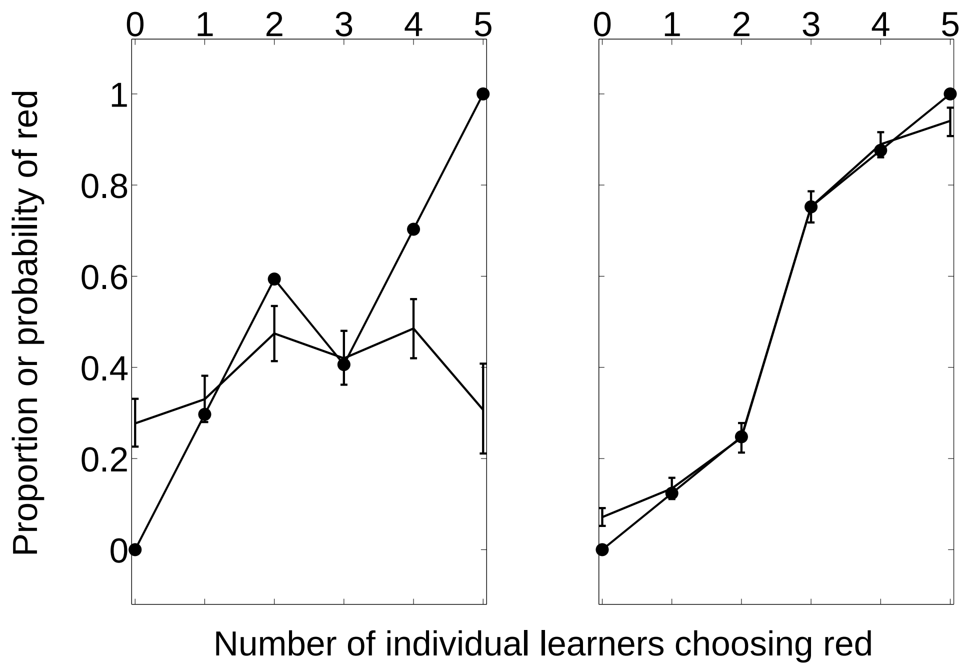

Figure 2: Data and theoretical predictions for social learners who were not stated conformists (left-hand panel) and those who were (right-hand panel). In both cases the graphs show the proportion of social learners choosing red (lines with 95% bootstrapped confidence intervals) as a function of the number of individual learners choosing red. The graphs also show (lines with solid circles) the theoretical probability a social learner chooses red under the frequency-dependent model described in the text and the maximum likelihood estimate of D. The MLE estimate for the 12 players who were not stated conformists is -0.4843 and the standard error is 0.0438. For stated conformists the estimate and standard error are 0.3805 and 0.0250 respectively. The different point estimates of D for the two groups are responsible for the different theoretical predictions. See the electronic supplement for additional detailed results.

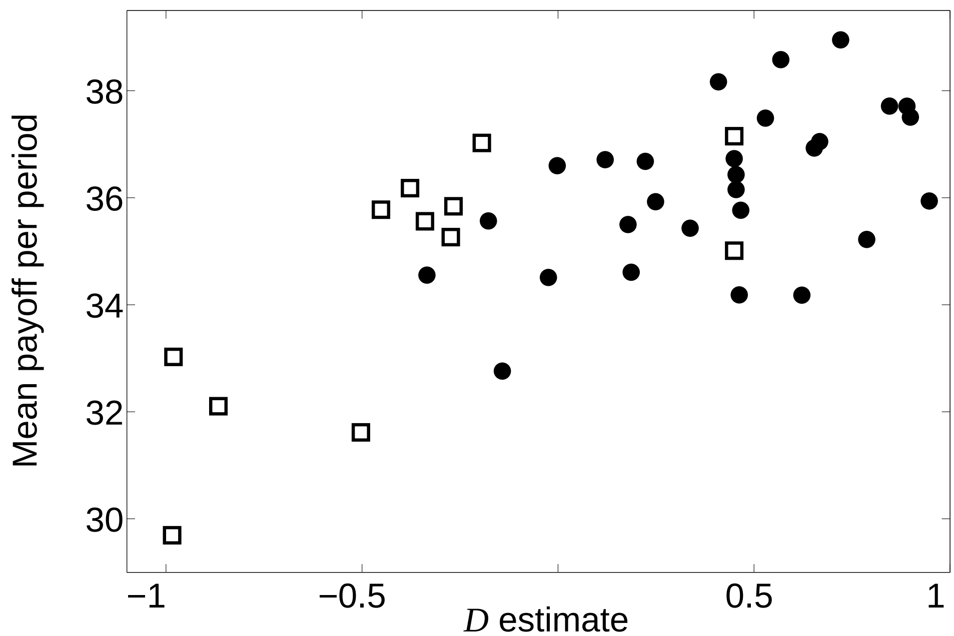

Figure 3: The mean payoff per period versus the estimated value of D for each of the 40 social learners. D estimates are based on the model of individual fixed effects. Players who were not stated conformists are shown with open squares, while stated conformists are shown in solid circles.

Experimental methods

Experiments were conducted in the laboratory of the Institute for Empirical Research in Economics at the University of Zürich and implemented entirely on a local computer network using the z-Tree software developed by Fischbacher (1999). We recruited a total of 70 undergraduate students from the University of Zürich and the Swiss Federal Institute of Technology. We ran a total of two sessions with 35 students in each session.

Students in each session were divided into 3 "worlds." Each world consisted of two groups. A group of 5 individuals received individual payoff information as described in detail below. These players were the *individual learners*. A group of 6 or 7 players who played simultaneously received no individual payoff information during the experiment. Instead, these *social learners* received information on the distribution of choices among the individual learners as described below. The only payoff information a social learner received was her total earnings after the experiment and questionnaire were finished.

Subjects were first informed they were participating in a laboratory experiment at the University of Zürich. Communication between subjects was not permitted. Students earned points in the experiment, and 150 points were worth one Swiss franc (about 0.78 USD or 0.64 EUR).

Subjects were instructed that their task was to choose either a "blue" technology or a "red" technology in each period. Technologies generated points at random according to specific probability distributions. One technology, the optimal technology, had a higher average payoff but was otherwise like the sub-optimal technology. The color of the optimal technology was chosen at random with probability 0.5.

The sub-optimal technology generated draws from a normal distribution with an expectation of 30 points and a standard deviation of 12 points. The optimal technology had a normal payoff distribution with an expectation of 38 points and a standard deviation of 12 points. Both distributions were truncated at 0 and 68. Truncation means that the truncated and untruncated distributions had slightly different expectations and standard deviations. Payoffs were rounded to integer values, and thus the set of possible payoffs for both colors was $\{0, 1, \ldots, 68\}$. We did not describe the payoff distributions to subjects in technical terms, but before they began the actual experiment we did provide them with an intuitively accessible demonstration of the random processes governing payoffs.

Specifically, the instructions before the experiment paid particular attention to the possibility that subjects may not have understood the formal concept of payoffs that follow probability distributions. Before the experiment started, subjects saw a demonstration of the two random technologies. In the demonstration, the optimal color was first determined randomly with probability 0.5. The optimal color was the same for each subject within a world but potentially different across worlds. Once an optimal color had been determined, the computer would take 250 draws from each of the two probability distributions for each subject individually. Two horizontal number lines from 0 to 68 appeared one beneath the other on the screen with one number line for each color. For each draw producing a specific value (e.g. 27 points for a red choice), the computer would place a little box (colored red or blue) along the appropriate number line (e.g. a red box at 27 for the number line being used to plot red draws). For multiple draws producing the same payoff, boxes were stacked on top of each other. As a consequence, students essentially watched a histogram being built draw by draw on the screen in front of them. This allowed them an intuitive sense of the stochastic process even if they had no training in probability theory or data analysis. Moreover, while the histograms were being built, they knew which color was optimal (but *only* for the demonstration). They could thus see, as an example, that blue was producing payoffs centered around 38, while red was producing payoffs centered around 30. The histogram was explained to them in writing, and subjects could read the explanation repeatedly while the histograms were being built. After the first demonstration, an optimal color was selected, and the demonstration was repeated. The whole demonstration phase took 5-10 minutes.

After the demonstration had been completed, subjects were informed that one repetition of the experiment would last for 25 periods. The timing was as follows. The computer first assigned the optimal color at random. The optimal color stayed the same for all 25 periods and was identical for each subject within a world. In each period, subjects chose between red and blue by indicating the desired color on the computer screen and clicking "OK."

Immediately after making a choice, individual learners were informed privately about the realized number of points received. Subjects also knew that the points from each period would be added up to yield a total payoff at the end of the entire experimental session.

Social learners were informed of the fact that in "the other part of the laboratory" a group of five individuals was facing the same two technologies with the same optimal color. Importantly, social learners also knew that subjects in the other group knew how many points they were making after each choice. In essence social learners knew that the players in the other group were receiving individual feedback.

Within each period, social learners were first informed about the number of individuals in the other group choosing red and the number choosing blue. Social learners then made their own choices. Social learners knew they would not receive information about their own payoffs until the entire session was over for the day.

The experiment was repeated six times. After the six repetitions, subjects responded to a questionnaire that recorded basic socio-demographic characteristics like gender, age, and academic major. The questionnaire also asked about learning strategies. Individual learners were questioned about the events that led them to revise their choices. Social learners were asked whether they tended to choose (i) the same color as the majority of the players in the other group, (ii) the same color as the minority of the players in the other group, or (iii) none of the two. The main paper labels social learners answering (i) as "stated conformists," while social learners answering (ii) or (iii) are labeled as "not stated conformists."

After the questionnaire, subjects received their total payoffs based on points summed over all six repetitions. The average payoff was 32.68 CHF (25.50 USD; 20.91 EUR). Sessions lasted about 2 hours.

How conformity works

In an environment with two behaviors or technologies, stochastic payoffs, and information about how common different behaviors are, a fundamental question centers around the ability of a frequency-dependent rule like conformity to filter noisy individual feedback into a useful social signal. As a simple approach to this problem, consider a group of N individual learners, and each of them chooses the best technology with probability $\overline{p}$. In this case the probability a majority of individual learners chooses this optimal technology when N is odd is simply

$$P(\text{majority opt}) = \sum_{i=\lceil N/2 \rceil}^{N} \binom{N}{i} \overline{p}^{i} (1-\overline{p})^{N-i}. \quad (1)$$

Figure S1 shows how $P(\text{majority opt})$ varies as a function of $\overline{p}$ for 4 different values of N.

If $\overline{p}$ is the probability each individual learner chooses the optimal technology, then $P(\text{majority opt}) < \overline{p}$ if $\overline{p} < 0.5$, but $P(\text{majority opt}) > \overline{p}$ if $\overline{p} > 0.5$. This fact is the essence of conformity's power to reduce noisy individual feedback into a useful social signal. Conformity works by identifying the optimum disproportionately if other forces, as summarized by $\overline{p}$, bias choices toward the optimum. It does not work, in the sense that it disproportionately identifies the sub-optimal technology, if other forces bias choices toward

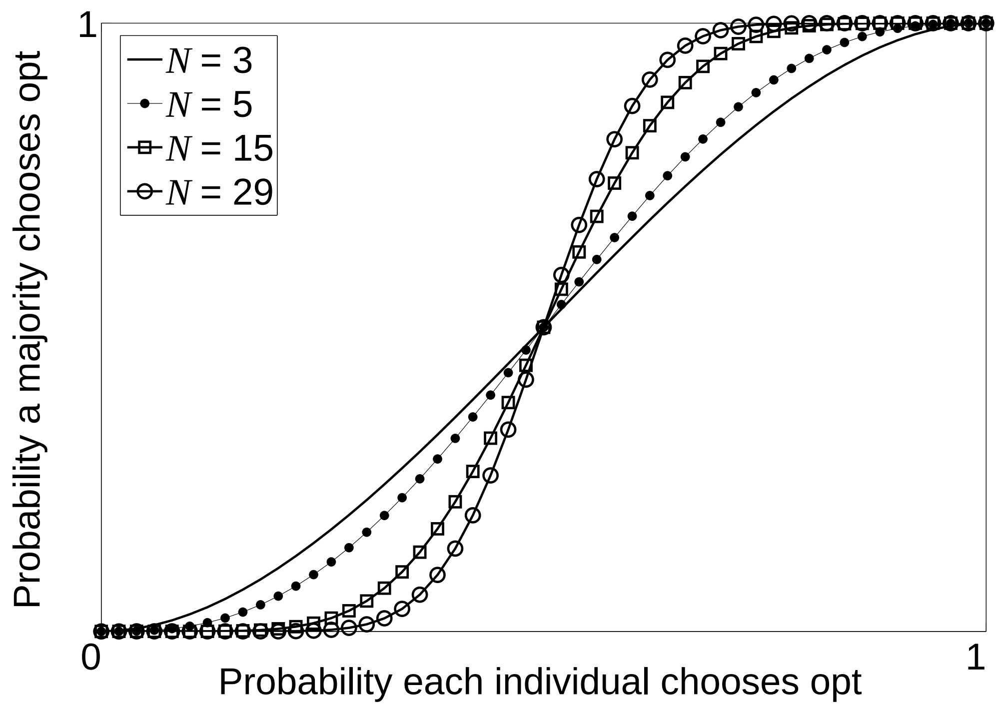

Figure S1. The probability a majority of individuals in the group, $P(\text{majority opt})$, chooses the optimal technology as a function of the probability, $\overline{p}$, that each individual chooses the optimal technology for 4 different group sizes, N.

the sub-optimum. Moreover, as the social group gets larger, the amount of information embedded in the group increases, and the nonlinearity intrinsic to conformity becomes more extreme. In sum, conformity exaggerates the effectiveness of individual learning. If individual learning is bad, conformity is worse. If individual learning is good, conformity is even better. Non-conformity, which we discuss briefly below, has exactly the opposite properties.

To expand this theoretical discussion, let us work with a simplified model that pertains more closely to the experiment. The two technologies with random payoffs are "red" and "blue." Assume two groups, individual learners and social learners, as in the experiment. Posit a focal social learner who assumes every individual learner in period t chooses red, given that red is optimal, with probability p_t. The social learner further assumes every individual learner chooses red, given that blue is optimal, with probability v_t. One particularly simple form of conformity ignores how overwhelming the majority is in a given period and simply notes which color was chosen by a majority of individual learners. In the present experiment, individual learners were always in groups of size 5. How does noting the color in the majority among individual learners provide useful social information?

To answer this question, note that the probability a majority of individual learners chooses red (maj red) given that red is optimal (red opt), which we call $f(p_t)$, is

$$P(\text{maj red} \mid \text{red opt}) = \sum_{i=3}^{5} \binom{5}{i} (p_t)^i (1 - p_t)^{5-i} = f(p_t). \tag{2}$$

The probability a minority chooses red (min red) given red is optimal is

$$P(\text{min red} \mid \text{red opt}) = \sum_{i=0}^{2} \binom{5}{i} (p_t)^i (1 - p_t)^{5-i} = 1 - f(p_t). \tag{3}$$

Similarly, the probability a majority of individual learners chooses red given blue is optimal (blue opt), which we call $h(v_t)$, is

$$P(\text{maj red} \mid \text{blue opt}) = \sum_{i=3}^{5} \binom{5}{i} (v_t)^i (1 - v_t)^{5-i} = h(v_t), \tag{4}$$

while the analogous probability red is in the minority is

$$P(\text{min red} \mid \text{blue opt}) = \sum_{i=0}^{2} \binom{5}{i} (v_t)^i (1 - v_t)^{5-i} = 1 - h(v_t). \tag{5}$$

Define s_t as the prior probability for a social learner that red is optimal (i.e. the probability that applies before the social learner knows the distribution of choices among individual learners in t). Bayes' rule specifies

$$P(\text{red opt} \mid \text{maj red}) = \frac{f(p_t)s_t}{f(p_t)s_t + h(v_t)(1 - s_t)}, \tag{6}$$

and

$$P(\text{red opt} \mid \text{min red}) = \frac{(1 - f(p_t))s_t}{(1 - f(p_t))s_t + (1 - h(v_t))(1 - s_t)}. \tag{7}$$

To simplify matters further, assume the focal social learner believes that $p_t = 1 - v_t \Rightarrow v_t = 1 - p_t$. This assumption simply means the social learner believes individual learners have no color biases in that they are equally likely to choose the optimal color regardless of whether it is red or blue. Specifically it means that p_t is the assumed probability an individual learner chooses red given red is optimal, the assumed probability an individual learner chooses blue given blue is optimal, and thus simply the assumed probability an individual learner chooses the optimal color. In turn, it means $1 - p_t$ is the assumed probability an individual learner chooses blue given red is optimal, the assumed probability an individual learner chooses red given blue is optimal, and thus simply the assumed probability an individual learner chooses the sub-optimal color.

Correspondingly, if blue is the optimal technology (blue opt), we can define $g(p_t)$ by rewriting equations 4 and 5 in terms of p_t,

$$P(\text{maj red} \mid \text{blue opt}) = \sum_{i=3}^{5} \binom{5}{i} (1-p_t)^i (p_t)^{5-i} = g(p_t) = h(v_t), \tag{8}$$

and the probability of a red minority is

$$P(\text{min red} \mid \text{blue opt}) = \sum_{i=0}^{2} \binom{5}{i} (1-p_t)^i (p_t)^{5-i} = 1 - g(p_t) = 1 - h(v_t). \tag{9}$$

Further define $m_t \in \{0,1\}$, where $m_t = 0$ means an observed minority of the individual learners chose red in t, while $m_t = 1$ means an observed majority chose red. The updated probability that red is optimal for a given social learner is then

$$s_{t+1} = \frac{m_t f(p_t) s_t}{f(p_t) s_t + g(p_t)(1-s_t)} + \frac{(1-m_t)(1-f(p_t)) s_t}{(1-f(p_t)) s_t + (1-g(p_t))(1-s_t)}. \tag{10}$$

The updating equation (10) tells us that only one of the two conditional probabilities, (6) or (7), is relevant in any given period, but which one is relevant depends on whether the social learner observed a majority or minority of red choices among the individual learners. The unconditional updated probability, s_{t+1} is a function of s_t, of course, but it also depends on p_t, a quantity that captures what the social learner thinks about how individual learners are learning. Figures S2 and S3 show the probabilities specified by (6) and (7) for five different values of p_t.

Specifically, the figures offer a simple formalization of the interplay between individual learning and conformity for social learners in the present experiment. After observing whether red or blue was in the majority among individual learners, the Bayesian social learner we have modeled should choose the color with an unconditional updated probability

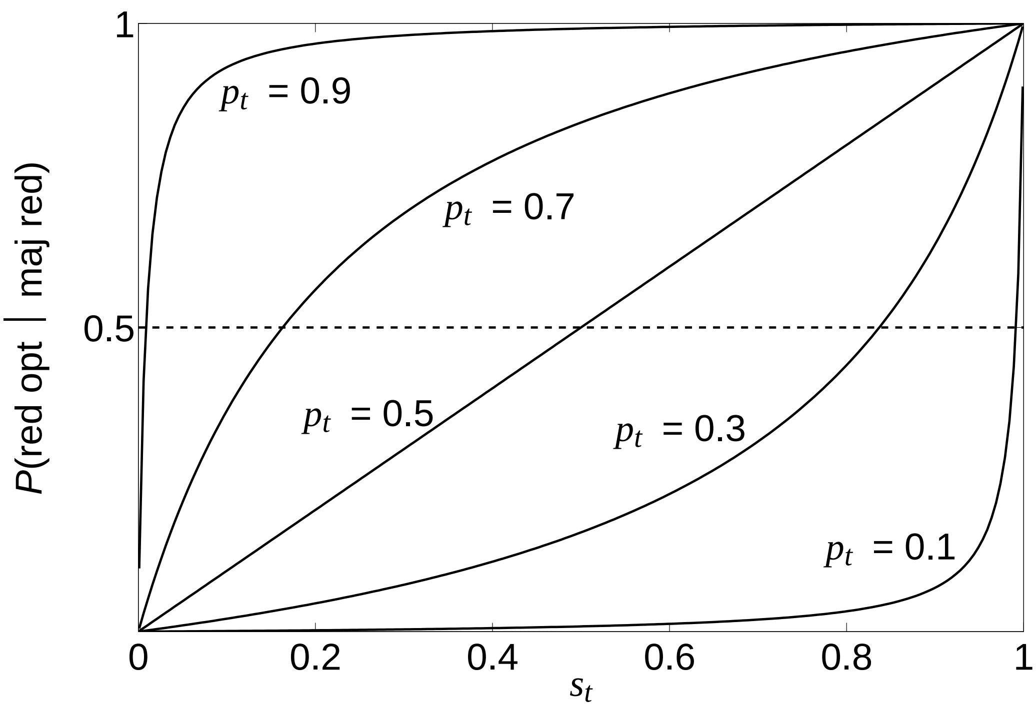

Figure S2. The updated probability that red is the optimal technology, given that a majority of individual learners chose red, as a function of s_t, the lagged unconditional probability that red is optimal. The function is shown for five different values of p_t, the social learners belief about the probability that each individual learner is choosing optimally in t.

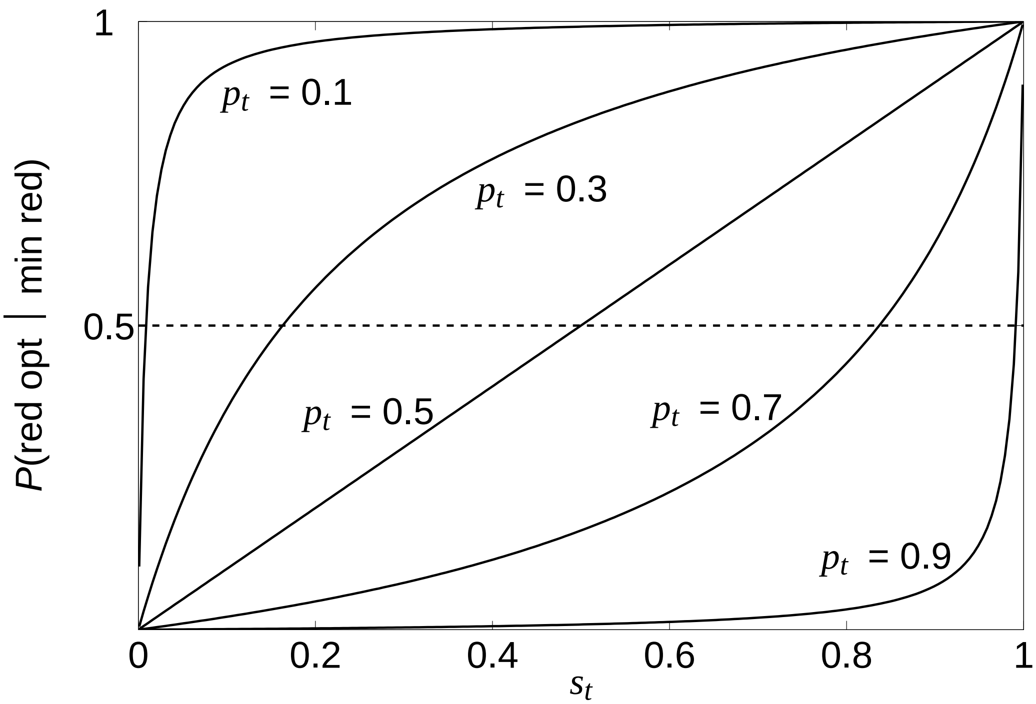

Figure S3. The updated probability that red is the optimal technology, given that a minority of individual learners chose red, as a function of s_t, the lagged unconditional probability that red is optimal. The function is shown for five different values of p_t, the social learners belief about the probability that each individual learner is choosing optimally in t.

of being the optimum that is greater than 0.5. This updated probability depends on the observed distribution of choices among the individual learners (specifically on whether red was in the majority or minority), the social learner's prior unconditional beliefs about the probability that red is optimal (s_t), and the social learner's beliefs about the decision-making of the individual learners as summarized by p_t.

As the figures show, if the social learner believes that individual learners are flipping a coin to make choices (i.e. $p_t = 0.5$), the information about the distribution of choices among individual learners is irrelevant. Updating essentially does not occur because the social learner's posterior probability, s_{t+1}, is equal to her prior probability, s_t. If, however, the social learner believes individual learners are biasing their choices in some way (i.e. $p_t \neq 0.5$), the social learner can use the information about the distribution of behaviors among individual learners to bias her own choice toward the optimum. How she uses the social information depends on her beliefs about the biases exhibited by individual learners.

Specifically, if the social learner believes individual learners are biasing their choices toward the optimum (i.e. $p_t > 0.5$), updating will tend to produce new beliefs that favor selecting the color in the majority among individual learners. This effect is stronger as p_t increases, s_t increases, or both. As the Figures S2 and S3 show, if p_t or s_t is close to 0 or 1, the variable with an extreme value has an overwhelming effect on updating unless the other variable has an equally extreme value with a countervailing effect.

One approach to updating would be to posit that social learners are initially neutral about which color is optimal (i.e. $s_0 = 0.5$), and they also believe individual learners are initially neutral ($p_0 = 0.5$). If social learners additionally believe individual learning is effective in the sense that $\forall t$, $p_{t+1} \geq p_t$, where the inequality is strict for some t, then these assumptions will have various implications for how social learning proceeds. In particular, as p_t increases, the updating rule implies that social learners should become more responsive to social information; the tendency to conform should become stronger through time. Moreover, if social learners vary notably with respect to p_t in any given period, some will be more responsive to changes in the social signal than others. For example, consider a situation in which red has been in the majority among individual learners. In period $t+1$, two social learners have a prior of $s_t = 0.9$, but one believes $p_t = 0.7$, while the other believes $p_t = 0.9$. Assume that, in contrast to period $t-1$, red is in the minority among individual learners in t. According to Figure S3, the social learner who believes $p_t = 0.7$ will have an updated belief that satisfies $s_{t+1} > 0.5$, while the social learner who believes $p_t = 0.9$ will update such that $s_{t+1} < 0.5$. The former will still choose red, while the

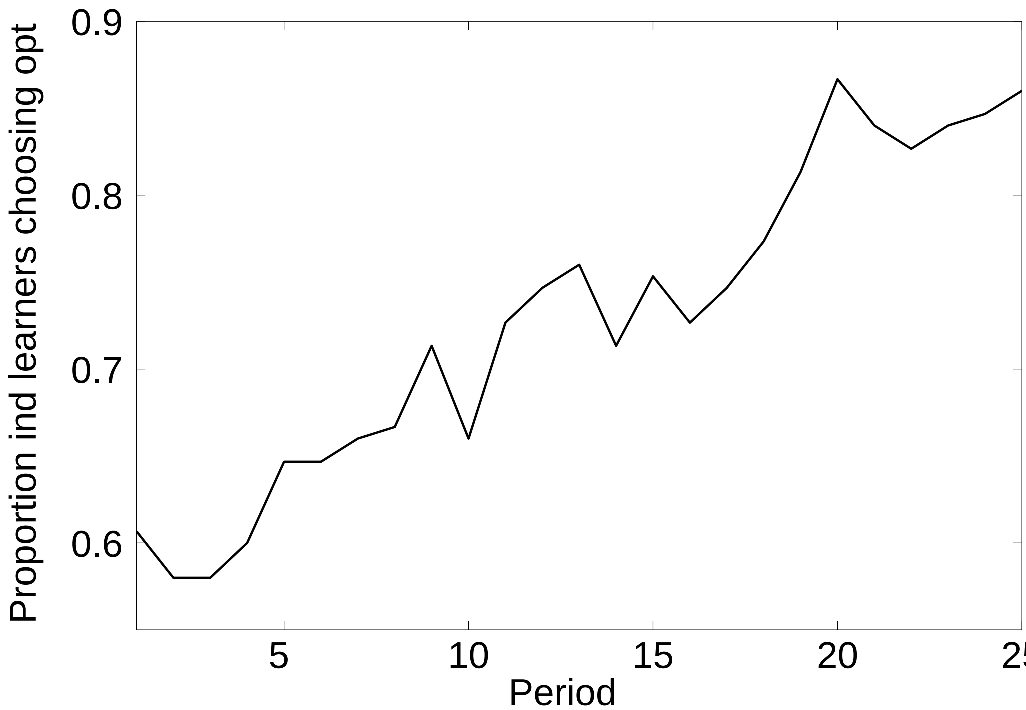

Figure S4. The proportion of individual learners choosing the optimal technology through time. The upward trend is highly significant ($p < 0.01$) when we regress the proportion choosing optimally on period using the method of Newey and West (1987) to correct for heteroskedasticity and autocorrelation up to lag 3. In this case, the estimated coefficient for period is 0.012 and the R^2 value is 0.930.

latter will switch to blue. Such variation could be one source of noise in the data for social learners. In general we do not claim that social learners in our experiment were Bayesians, but the updating model we have presented provides a convenient approach to summarizing the value of conformity as it interacts with individual learning. Figure S4 shows that the individual learners in the experiment were learning effectively.

Classic theory

Boyd and Richerson (1982, 1985) developed a simple model of conformity that has the following properties. In $t+1$ each individual in the population samples N individuals from the previous period. To simplify matters in an unimportant way, assume N is an odd number. Define $I_t \in \{0, 1, \ldots, N\}$ as a random variable with realizations i_t specifying the number of individuals choosing red in a particular sample of size N. At the individual level conformity takes the following form,

$$\begin{aligned} i_t < N/2 \quad &\Rightarrow \quad P(\text{focal ind chooses red}) = \frac{i_t}{N}(1-D) \qquad (11) \\ i_t > N/2 \quad &\Rightarrow \quad P(\text{focal ind chooses red}) = \frac{i_t}{N} + \frac{D(N-i_t)}{N}, \end{aligned}$$

where $D \in [-1, 1]$ is a parameter that measures the strength of conformity. Taking account of the fact that each individual samples from the population, and thus by chance all individuals do not have the same information about the distribution of behaviors, dynamics in a large population change according to

$$r_{t+1} = \frac{1-D}{N} + D \sum_{i_t=\lceil N/2 \rceil}^{N} \binom{N}{i_t} (r_t)^{i_t} (1-r_t)^{N-i_t}. \qquad (12)$$

In a model combining individual and social learning, Boyd and Richerson (1989) determined that the evolutionarily stable value of D is 1. When $D = 1$, equation (12) is equivalent to (1). Consequently, in addition to our earlier discussion, we now see that Figure S1 also shows a set of functions that specify the recursions for the dynamics of a population of $D = 1$ conformists. The difference is one of interpretation. Under the population dynamics interpretation, we assume the horizontal axis in Figure S1 is r_t rather than $\overline{p}$. We also assume that each individual *samples* N others from a large population. Given this interpretation, as N increases, the nonlinearity intrinsic to conformity increases because all individuals have more accurate estimates of the real distribution of behaviors in the entire population. Everyone uses social information in exactly the same way (i.e. $D = 1$), and their estimates of the distribution of behaviors become increasingly correlated as N increases.

If $D = 1$ and $N = 5$, we have a simple theoretical prediction for the social learners in our experiment,

$$P(\text{soc learners chooses red}) = \sum_{i_t=3}^{5} \binom{5}{i_t} (r_t)^{i_t} (1 - r_t)^{5-i_t}. \qquad (13)$$

Figure S5 shows the match between this theoretical prediction and the data for the subsample we describe in the main paper as "stated conformists." The match is relatively good, but it is good for the wrong reason. Model (12) includes two sources of randomness: (1) individuals exhibit an incomplete propensity to conform as measured by D, and (2) individuals differ in their assessment of which behavior is in the majority because they sample from the population. To derive the prediction specified by (13), we removed one source of randomness by setting $D = 1$. The only remaining source of noise in the theoretical prediction is due to sampling. Nonetheless, in spite of the fit between the data and the prediction, sampling cannot be the source of noise in the data because we provided full distributional information exogenously in all periods of the experiment. In both cases noise is binomial, but this exercise highlights how the proposed mechanisms are very different. In one case, binomial noise exists because social learners have an incomplete tendency to conform. In the other case, they vary in terms of which behavior they think is in the majority because none of them have perfect information about the distribution of behaviors in the population. We removed the latter possibility via our experimental design, so the noise in the data must be due to an incomplete tendency to conform.

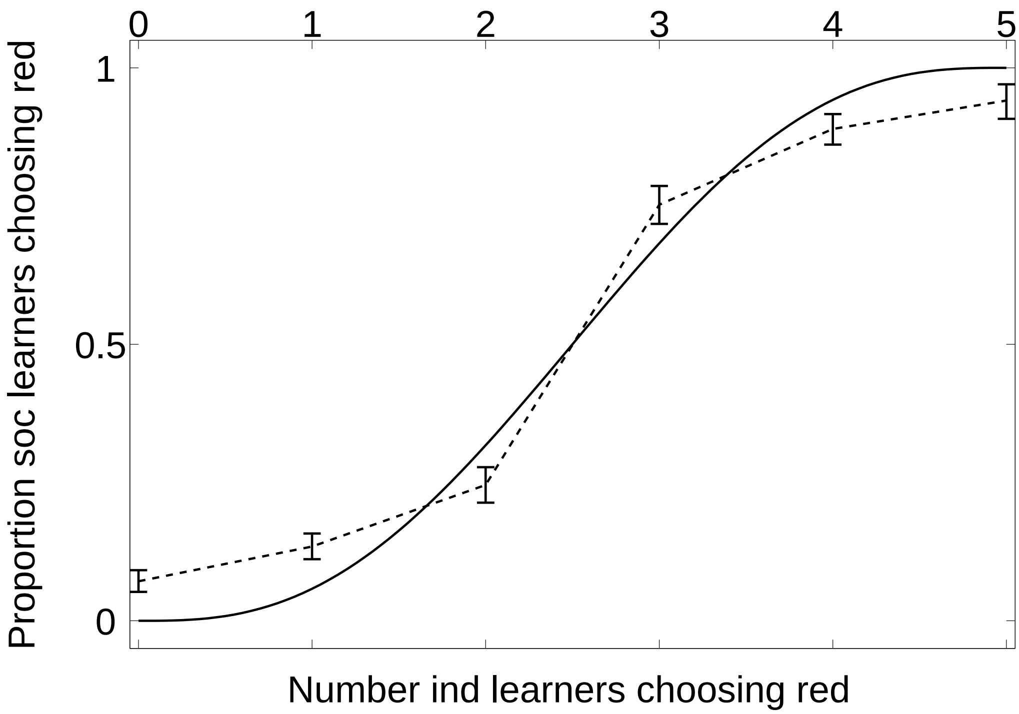

Figure S5. The theoretical probability or observed proportion of social learners choosing red versus r_t, the proportion of individual learners choosing red. The theoretical prediction (solid line) is based on the conformity model with $D = 1$ and the assumption that each individual *samples* $N = 5$ others from the population. The data shown (dashed line with error bars) are only for the sub-sample of social learners who claimed they followed the majority during the experiment ("stated conformists"), and 95% bootstrapped confidence intervals are included.

Data and the monotonicity problem

To clarify what an experiment must do to show conformity, treat model (12) as an individual probability model. Let $N = 3$ and consider three values of D. When $D = 1$, the model is a conformity model,

$$P(\text{soc learner chooses red}) = r_t + r_t(1 - r_t)(2r_t - 1). \tag{14}$$

When $D = 0$, the model is simply a model of linear social learning as examined in Boyd and Richerson (1985),

$$P(\text{soc learner chooses red}) = r_t. \tag{15}$$

When $D = -1$, the model is non-conformist,

$$P(\text{soc learner chooses red}) = r_t + r_t(1 - r_t)(1 - 2r_t). \tag{16}$$

In all three cases, as Figure S6 shows, the probability a social learner chooses red increases monotonically with r_t. Consequently, a researcher cannot argue conformity by simply showing that individuals are more likely to adopt a given behavior as it becomes more common. All three of these models make that prediction. Nonetheless, distinguishing among them is critical because at the aggregate level they have profoundly different dynamics and correspondingly make profoundly different predictions. To identify conformity cleanly, researchers must identify a particular kind of nonlinear, monotonically increasing relationship between current and future behaviors. Equally important, as suggested in the previous section, the researcher must account for whether social learners do or do not sample behavioral models from a larger population as this distinction affects the prediction under different frequency-dependent heuristics.

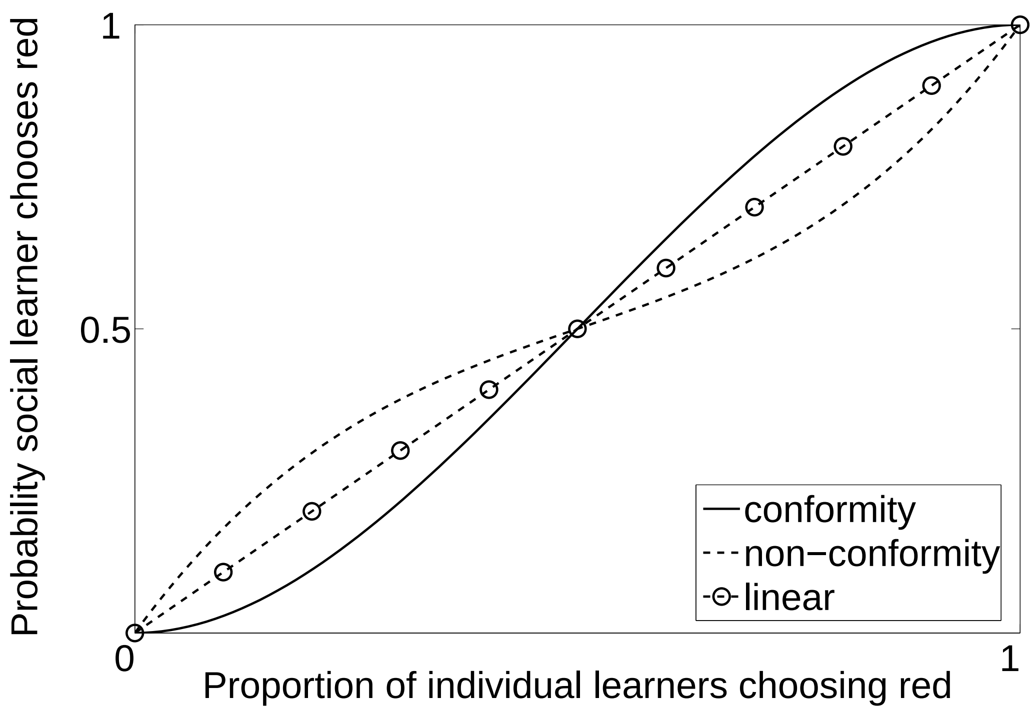

Figure S6. The probability a social learner chooses red as a function of r_t, the proportion of individual learners choosing red. Models shown are conformity, linear social learning, and non-conformity when $D = 1$ and each individual samples $N = 3$ others from the population.

Statistical models and detailed results

To estimate conformity using (11) and data from the present experiment, let $k \in K = \{1, 2, \ldots, BT\}$, where K is a set to index observations for each social learner by period combination. B is the number of social learners, and T the number of periods. Define $c_k \in \{0, 1\}$ as a variable that records the choice such that $c_k = 0$ if the social learner chose blue, and $c_k = 1$ if the social learner chose red. Call the entire data set $c = \{0, 1\}^{BT}$. Further define $i_k \in \{0, 1, \ldots, 5\}$ as an associated variable that records the number of individual learners choosing red as observed by the social learner in question in the appropriate period. The probability model for a given social learner is thus

$$\begin{aligned} i_k \leq 2 \quad &\Rightarrow \quad P(c_k = 1) = \frac{i_k}{5}(1 - D) \\ i_k \geq 3 \quad &\Rightarrow \quad P(c_k = 1) = \frac{i_k}{5} + \frac{D(5 - i_k)}{5}. \end{aligned} \tag{17}$$

To estimate D, one must remove all data points where $i_k \in \{0, 5\}$ because observations in which all individual learners chose the same option cannot be used to estimate D. This results simply from the fact that D drops out of model (11) in these cases. Accordingly, define the set $J = \{k \in K \mid i_k \in \{1, 2, 3, 4\}\}$. The log likelihood function for estimating D is thus

$$\ln\{L(D \mid c)\} = \sum_{k \in J} c_k \ln\left\{\frac{P(c_k = 1)}{P(c_k = 0)}\right\} + \ln\{P(c_k = 0)\}. \tag{18}$$

This function is derived from a joint probability distribution over c given that observations are Bernoulli random variables under (17). We used this method to estimate the values of D reported in the main text, and approximate standard errors were calculated by inverting the Hessian of the log likelihood function evaluated at the estimated value of D.

Incorporating heterogeneity into the model (e.g. "not stated conformists" versus "stated

Table S1. Estimates, standard errors, and approximate 95% confidence intervals for the model that assumes a single D value for all social learners and the model that estimates a separate D for the social learners who were stated conformists (D_Y) and those who were not (D_N). Confidence intervals are the point estimates plus or minus twice the standard errors.

Model	Parameter	Estimate	Std. error	95% CI
Single D	D	0.1081	0.0005	$[0.1070, 0.1091]$
Conformist (Y or N)	D_Y	0.3805	0.0250	$[0.3305, 0.4305]$
	D_N	-0.4843	0.0438	$[-0.5720, -0.3966]$

conformists") is equivalent to splitting the sample in an appropriate way, fitting the model separately to each sub-sample, and then adding maximized log likelihood values to calculate Akaike values for the entire data set. In practice, however, one can write a routine that accommodates any desired degree of heterogeneity. Code for implementing these procedures in R (R Development Core Team, 2006) is available on request. Apart from the model-fitting results presented in the main text, Tables S1 and S2 show additional detailed results from each of the three models.

Table S2. Estimates, standard errors, and approximate 95% confidence intervals for the fixed effects model that includes a separate D for each social learner. Confidence intervals are the point estimates plus or minus twice the standard errors.

Parameter	Estimate	Std. error	95% CI
D_1	0.4542	0.1193	$[0.2155, 0.6928]$
D_2	0.6538	0.0979	$[0.4580, 0.8497]$
D_3	0.4622	0.1176	$[0.2269, 0.6975]$
D_4	-0.2736	0.1489	$[-0.5713, 0.0241]$
D_5	-0.4517	0.1404	$[-0.7325, -0.1708]$
D_6	0.1865	0.1347	$[-0.0829, 0.4559]$
D_7	-0.5028	0.1391	$[-0.7809, -0.2246]$
D_8	-0.3394	0.1550	$[-0.6495, -0.0294]$
D_9	-0.0020	0.1553	$[-0.3127, 0.3087]$
D_{10}	-0.1944	0.1607	$[-0.5158, 0.1270]$
D_{11}	-0.8665	0.1419	$[-1.1503, -0.5827]$
D_{12}	0.8456	0.0752	$[0.6953, 0.9960]$
D_{13}	0.5291	0.1249	$[0.2792, 0.7791]$
D_{14}	0.1202	0.1580	$[-0.1959, 0.4363]$
D_{15}	-0.3344	0.1458	$[-0.6260, -0.0427]$
D_{16}	-0.2667	0.1447	$[-0.5561, 0.0226]$
D_{17}	0.4542	0.1166	$[0.2210, 0.6874]$
D_{18}	0.5683	0.1071	$[0.3541, 0.7824]$
D_{19}	-0.9809	0.1083	$[-1.1977, -0.7642]$
D_{20}	0.4093	0.1220	$[0.1653, 0.6533]$

Continued on next page

Parameter	Estimate	Std. error	95% CI
D_{21}	-0.1777	0.1435	$[-0.4646, 0.1093]$
D_{22}	0.3370	0.1701	$[-0.0033, 0.6773]$
D_{23}	0.1787	0.1819	$[-0.1851, 0.5425]$
D_{24}	0.8902	0.0763	$[0.7377, 1.0428]$
D_{25}	0.6675	0.1285	$[0.4104, 0.9245]$
D_{26}	0.7211	0.1192	$[0.4826, 0.9596]$
D_{27}	0.2226	0.1726	$[-0.1225, 0.5677]$
D_{28}	-0.3775	0.1853	$[-0.7480, -0.0069]$
D_{29}	0.4493	0.1589	$[0.1314, 0.7673]$
D_{30}	-0.0246	0.1931	$[-0.4108, 0.3617]$
D_{31}	0.9471	0.0524	$[0.8422, 1.0520]$
D_{32}	0.4493	0.1589	$[0.1314, 0.7673]$
D_{33}	0.7875	0.1027	$[0.5822, 0.9929]$
D_{34}	0.4493	0.1589	$[0.1314, 0.7673]$
D_{35}	-0.9844	0.0933	$[-1.1710, -0.7978]$
D_{36}	0.4663	0.1049	$[0.2566, 0.6760]$
D_{37}	0.6220	0.0908	$[0.4403, 0.8036]$
D_{38}	0.2489	0.1196	$[0.0097, 0.4881]$
D_{39}	0.8990	0.0496	$[0.7998, 0.9982]$
D_{40}	-0.1419	0.1139	$[-0.3697, 0.0857]$

References

Boyd, R. and Richerson, P. J. (1982). Cultural transmission and the evolution of cooperative behavior. *Human Ecology*, **10**(3), 325–351.

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

Boyd, R. and Richerson, P. J. (1989). Social learning as an adaptation. In A. Hastings, editor, *Lectures on Mathematics in the Life Sciences*, pages 20: 1–26. American Mathematical Society.

Fischbacher, U. (1999). z-Tree: toolbox for readymade economic experiments. IEW Working paper 21, University of Zuerich.

Newey, W. K. and West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, **55**(3), 703–708.

R Development Core Team (2006). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0.